

The Job Ad Guide

For data professionals who want to understand what they are actually
applying for

Lauri Koobas

2026

Table of Contents

Why it matters	3
How to use this guide	3
The Core Framework: Capacity vs Capability	3
The Eight Types of Job Ads	4
Type A: Capacity in a Mature Team	4
Type B: Capacity in a New or Small Team	4
Type C: Capability Disguised as Capacity	5
Type D: Capability – New Function in an Existing Team	5
Type E: Capability – Greenfield	6
Type F: Capability – The Impossible Job (Chaos & Unicorns)	6
Type G: Team Lead	7
Type H: Head of Data/Analytics	8
Depth vs Breadth	8
Reading the Technical Requirements	9
The Generic List	9
Where the Real Signal Lives	9
When You Don't Know the Exact Tool	9
Skills: Easy vs Hard to Learn	9
Domain Knowledge	10
Questions to Ask	10
Always Ask	10
Type-Specific Questions	11
Meta: Things Beyond the Job Content	12
The 100% Qualification Trap	12
Onsite / Hybrid / Remote	12
Relocation	13
Compensation	13
The Perks Section	13
Recruitment Agencies	13
Additional Research	14
Check for Other Roles at the Same Company	14
Red Flags Summary	14
Summary	14
Annex: Real-World Example – Type Analysis, 1	16
Annex: Real-World Example – Type Analysis, 2	19
Annex: Real-World Example – Type Analysis, 3	21
Annex: Real-World Example – Type Analysis, 4	24

Why it matters

Most people read job ads like shopping lists. They scan for keywords, check if they match the requirements, and hit “Apply”. Or, as it sometimes seems from the hiring manager side - they skip all the other parts and just hit “Apply”. This is backwards.

Job ads are documents written by humans who may or may not understand the role they’re hiring for. Some are written by experienced hiring managers who know exactly what they need. Others are assembled by HR (human resources/people department) from a template, with bullet points contributed by three different people who have never talked to each other.

Learning to read between the lines will save you time and unearned disappointment. You’ll stop applying for roles that sound good but are actually chaos. You’ll start recognizing the ones where you would thrive. And when you do get an interview, you’ll know which questions to ask.

How to use this guide

I suggest reading through it, especially the real world examples at the end. But it also works if you put it into your favorite AI tool along with a job description and ask it to apply the framework found in this document.

The Core Framework: Capacity vs Capability

Before diving into specifics, there’s one distinction that explains most job ads.

Capacity means “we need more hands doing what we already do.” The team exists, the processes exist, the understanding exists. They just need another person. If you are starting out in the field of data, either out of school or switching careers, this is where you want to focus.

Capability means “we need to do something we can’t do yet.” This might be a new function, a new team, or replacing someone whose job nobody fully understood.

This distinction matters because it tells you what kind of work you will actually be doing. Capacity hires get to focus. Capability hires get to define things – which is either exciting or exhausting, depending on who you are. I have personally always been capability hire, but I myself have hired both kinds.

I have identified 8 types of job ads, described below. I have arrived at these types by having read at least a thousand job ads for data roles over the last decade, having worked in multiple different companies and roles, given public and private feedback on job ads on LinkedIn and having also written quite a few ones myself. Enjoy :)

The Eight Types of Job Ads

Type A: Capacity in a Mature Team

This is the most straightforward hire. An existing team with a mature understanding of what they need. The job description is consistent – one or two main focus areas, clear responsibilities, optional seniority level. It means that they are open to hiring a mid-level or a senior or sometimes even a junior into the same role.

How to spot it: The “what you would be doing” section reads like it was written by someone who actually does this job. The tasks make sense together. Nothing feels randomly added.

What it means: Backfill or growth. They know what they need because they’ve needed it before.

Good for: Almost everybody - because learning what a good mature setup and ways of working looks like is always useful. People who want stability. Anyone who prefers depth over breadth. If you’re tired of chaos and just want to do good work in a well-run environment, this is what you’re looking for.

Not the best fit for: People who want to live on the bleeding edge of technological advancement. You won’t be able to use programming languages, tools or AI models that are not approved.

Type B: Capacity in a New or Small Team

Still a capacity hire, but the team is young. Maybe two people, maybe just being formed. The job description is within a single domain (like data engineering), but the tasks are all over the place.

How to spot it: Lots of mentions of “setting up processes,” “establishing best practices,” and “making sure things run smoothly.” The responsibilities list is long and varied.

What it means: You’ll be building the function, not just working in it. Expect to wear many hats. The role will evolve as the team grows. Expect to enter into an environment with at least one large migration underway (read: hopelessly off track), either onprem to cloud or one platform to another, etc.

Good for: People who enjoy creating structure. Generalists who get bored doing one thing. Anyone who’s comfortable with ambiguity but still wants to work within a defined domain.

Not the best fit for: Someone who wants to just do a small set of things, like only analytics or only build machine learning models or only write Terraform. Also not the best for people who really don't want to make long-lasting architectural decisions - there might not be anyone there to help you make them.

Type C: Capability Disguised as Capacity

This one is tricky. On the surface, it looks like a normal hire. But reading closely, something feels off.

How to spot it: The tasks are within a domain like data engineering, but the descriptions are vague. Phrases like "maintain infrastructure," "ensure uptime," and "follow best practices" appear without specifics. It's often a replacement hire – someone left, and the company isn't entirely sure what that person was doing.

What it means: The job description was written by people who don't fully understand the role. They think they're hiring for capacity, but they're actually hiring for capability – they need someone to figure out what the job should be.

Good for: Proceed with caution. This can work if you're comfortable with ambiguity AND good at educating people who think they know what they need but don't. If both of those aren't true, keep scrolling.

Not the best fit for: People who mostly just want to do technical work. This type of a role will involve quite a lot of explaining why things take time, mostly to people who previously have not really cared about the specifics.

Type D: Capability – New Function in an Existing Team

An existing team is adding a genuinely new capability. Think of a data team hiring their first ML engineer, or an analytics team adding someone focused on experimentation.

How to spot it: Specific business outcomes mentioned, but not much technical guidance on how to achieve them. The role is clearly new – not a replacement. These job ads tend to be pretty clear in their language of "help us get to the next level of" kind.

What it means: Real expansion. You'll be defining what this function looks like at this company. The team exists to support you, but the specifics are yours to figure out.

Good for: Self-starters who can work without detailed instructions. People who enjoy having a clear business goal but freedom in execution. If you need someone to tell you what to do every day, this isn't it. Ideal for freelancers looking for a permanent position.

Not the best fit for: People who have not done this before in "full stack" as setting up the tools is going to be part of it. Also not ideal for someone who only wants to do the technical part as communicating with business counterparts is unavoidable.

Type E: Capability – Greenfield

The most extreme capability hire. You're not joining a team – you might BE the team. At least for now.

How to spot it: Phrases like "lead development of our data capability" or "help us become data-driven." The role is framed as building something from nothing.

What it means: You ARE the team. High autonomy, high ownership, high risk. Everything that goes right is yours. Everything that goes wrong is also yours. Note that this is different from "Head of" or "team lead" roles - this is an individual contributor role in a random team that might or might not evolve into something bigger over time.

Good for: High risk tolerance. People who genuinely want to build something from zero. You should have done this before (at least partially) or be very confident in your ability to figure it out. Also, make sure they're giving you budget and access to leadership – without those, you're set up to fail.

Not the best fit for: People with a heart condition, young family, or a lot of hobbies (or friends) they would like to keep. Working hours are for doing, evenings and weekends are for learning what to do.

This is how I became a data engineer - I joined into a tiny team, then my manager quit a few months later and I was alone in this role for the next three years.

Type F: Capability – The Impossible Job (Chaos & Unicorns)

The danger zone. Someone left – a one-man-band type who did a bit of everything. Nobody really understood what they did. Now the company needs to replace them.

How to spot it: The job description looks like it was assembled by asking three or four different people what that person did. The result is extremely disjointed. "Build data

strategy," "send weekly emails," "optimize cloud costs," and "update business glossary" all in one role. These tasks don't belong together.

What it means: Run. Or negotiate very hard. Expect to be interviewed by people with wildly different roles and have an essentially random manager.

Good for: If you think you're very good AND very bored, this could be fun. Or if you are really heavily leaning towards learning vs earning (or sanity), this might be a place to flex your intelligence, grit, and people skills.

Not the best fit for: Almost anyone. Really. I have interviewed a couple of times to such an ad just to make sure I'm not misreading the situation. I wasn't. One of those interviews turned into me helping them rewrite it into two separate ads. The recruiter was thankful.

Type G: Team Lead

Team lead roles vary wildly, and the job ad often doesn't make it clear what kind of a job it will be.

How to spot it: Count the bullet points in "what you would do." How many are hands-on technical work? How many are people management? This ratio tells you more than anything else in the ad. Note - the naming of the management roles are random, so check out the Head of X role below as well.

Watch out for: Compliance-related mentions. Things like "ensure processes are documented" or "maintain audit trails." This is neither technical work nor people management – it's change management. A whole different skill set that's rarely discussed in the interview.

What it means: The IC (individual contributor)/management split varies greatly between companies and even between teams. Some "team leads" write code 80% of the time. Others are in meetings all day. The job ad won't tell you which – you'll have to ask.

Good for: Depends entirely on what you want. If you want to stay technical, look for ads heavy on technical bullets. If you want to grow as a manager, look for the opposite.

Not the best fit for: People who are only passionate about the data-related aspects of the job. Doing one-on-ones, hiring and performance reviews will not leave you enough time to do the things you love doing.

Type H: Head of Data/Analytics

These roles are bimodal. Two completely different jobs share the same title.

Builder mode: The company is hiring a head to build a function from scratch. You'll be hiring your team, choosing your tools, setting your strategy. Exciting if you want to build.

Fixer mode: The company already has a data function, but it's dysfunctional. Brain drain, fires everywhere, nobody trusts the numbers. They need someone to stabilize things while also demonstrating value. Much harder than building, and often not acknowledged as such.

How to spot it: You often can't from the job ad alone. Both modes use similar language. The only way to know is to ask directly in the interview.

What it means: Ask: "Is this a build or a fix situation?" Listen carefully. If they hesitate or say "a bit of both," probe deeper.

Good for: Make sure you know which mode you're signing up for. They require different skills and offer very different experiences.

Not the best fit for: People who want to continue directly working with data. Neither of these modes contain much work as an individual contributor.

Depth vs Breadth

Separate from the capacity/capability distinction is the question of specialization.

Some roles are deeply specialized. "Marketing Mix Modeling expert." "Clickhouse DBA." "Experimentation Platform Engineer." These require specific experience that takes years to build. You either have it or you don't.

Other roles are broader. "Data Engineer." "Analytics Lead." "Senior Data Scientist." These assume general competence across a wider area. The specific tools and domains can be learned on the job.

Why it matters: Don't apply for deep specialization roles hoping to learn on the job. You won't get the interview. Do apply for broad roles even if you don't know their exact stack – competent generalists can pick up new tools.

Reading the Technical Requirements

The Generic List

Most job ads include a generic requirements section. "Cloud computing (GCP/AWS/Azure)." "Relational databases." "Data modeling experience."

This tells you almost nothing. Every company writes these same bullets. It's boilerplate.

Where the Real Signal Lives

The actual stack is hidden in the "nice to have" section. "Bonus points if you have..." "Would be great if you know..." "Experience with X is a plus."

This is where they reveal what they're actually using. If every "bonus" item is Databricks-related, they're a Databricks shop. If they mention dbt and Looker, that's probably the stack.

When You Don't Know the Exact Tool

Don't panic. Read up on it enough to have a conversation. Then talk about comparative experience.

"I haven't used PowerBI, but I've spent three years in Tableau. From what I understand, the core concepts translate – it would take me maybe a month to get productive."

Hiring managers know that specific tool experience is less important than fundamental skills. Show that you understand this, and the gap stops being a problem.

Skills: Easy vs Hard to Learn

Some skills transfer quickly. Learning a new BI tool when you know another one – a few weeks. Picking up a new cloud provider's managed services – a month or two.

Other skills take years. Statistical thinking. System design. Understanding how businesses actually work. These aren't things you cram before an interview.

Know the difference. Be honest about which gaps you can close quickly and which ones would take serious time.

This does not mean you should not apply - applying for curiosity is always OK. Besides, there is always a chance that this is the third round of candidates and everyone else before you has declined or not been suitable. This is how I got my first position as a data scientist in Microsoft!

Domain Knowledge

In some roles, domain knowledge is critical.

Marketing analytics requires understanding attribution models, customer acquisition costs, funnel thinking, and the psychology of how companies measure marketing spend. You can't fake this in an interview.

Science-adjacent roles (biotech, research) require understanding research methodology, publication processes, and the norms of academic collaboration.

Geo/spatial roles require knowing about coordinate systems, projections, and mapping conventions.

In other roles, domain knowledge is secondary. A data engineer at a fintech company doesn't need to understand financial products deeply – they need to build reliable pipelines. The domain context can be learned in the first few months.

How to spot which one you're dealing with: If the job ad spends two or more bullets on industry-specific context, domain matters. If it's all technical requirements with "experience in our industry is a plus" at the end, domain is learnable on the job.

Note: you don't need to fake an interest in the domain. It's enough to do the basics - read up on what the company does to the level where you can explain it to your parents. Then say that you did that during the interview and it's fine. Most companies care more about you being good at technical things.

Questions to Ask

Every interview should include questions from you. Here are some that go beyond the generic "tell me about the culture."

Always Ask

- "Why is this role open?"

- "What will be my first project?"
- "What does success look like in six months?"
- "What's the biggest challenge the team is facing right now?"

Type-Specific Questions

For Capacity in a Mature Team (Type A):

- "How did the last new person onboarding go?"
- "When can I start?"
- "Where do I sign?" :)

For Capacity in a New Team (Type B):

- "What are some long-term principles that the team is following?"
- "How does adding new tools work if needed?"

For Capability Disguised roles (Type C):

- "What happened to the previous person in this role?"
- "Where does work come from?"
- "Is there a single person who sets the priorities?"

For New Function roles (Type D):

- "Who is the main business stakeholder?"
- "Out of the three things that all together cannot be fixed at once - scope, time, and resources - which one is most flexible?"

For Greenfield roles (Type E):

- "What's the process for getting new tools approved and bought?"
- "Will there be a headcount for more people in this area soon?"

For Chaos Backfill roles (Type F):

- "What are the top three fires I'd need to put out in week one?"
- "Out of all the things mentioned in the job ad, are there people in other teams that could handle some of it?"

For Team Lead roles (Type G):

- "What has led the team to be in this particular part of the organization?"

- "Who and in which roles would be my main counterparts on a weekly basis?"
- "Is there internal training on how to be a people lead in this organization?"

For Head of Data roles (Type H):

- "Is the first priority to have less bad things or more good things? Putting out fires or getting more value out of it?"
- "What are some skeletons that I would discover in the first few months on the job?"

Meta: Things Beyond the Job Content

The 100% Qualification Trap

If you qualify for 100% of the requirements, you are overqualified.

All job ads are aspirational. They describe the ideal candidate – someone who might not exist. The actual hire will meet maybe 70-80% of the requirements.

The perfect example here is "technical degree", which has been a number one requirement for all of the jobs I have had. And I don't have a technical degree - mine is in psychology. It's never been a problem!

If you're at 70% and the role excites you, apply. You're probably right for it.

Onsite / Hybrid / Remote

Job ads sometimes say one thing and mean another.

"Hybrid" often means "we'd prefer you in office but legally can't require it." Ask what a hybrid actually looks like in practice for that particular team. Two days a week? Whenever there's a meeting? Officially flexible but culturally frowned upon?

"Remote-first" and "remote-friendly" are very different. Remote-first means the company operates as if everyone is remote. Remote-friendly means you can work remotely, but you'll miss out on things.

One very specific thing is usually non-negotiable - "Remote within US" or "Remote within EU". It's some legal requirement around them not being able to actually have an employment contract outside of a certain area. If you are not there, don't apply.

"Flexible" is meaningless until defined. Ask.

Relocation

If the role involves moving to a new city or country, factor that into your evaluation. A 20% salary bump doesn't mean much if the cost of living is 40% higher. Relocation support varies wildly – some companies cover everything, others give you a one-time payment that barely covers the moving truck.

Compensation

The EU Pay Transparency Directive is changing this landscape. In the near future, salary ranges will be required in job postings. Until then, you're often guessing.

Red flag:

- Early-stage startup that emphasizes equity: you will get paid very little real money. Equity in early-stage companies is usually worth nothing.

The Perks Section

Ignore all of it.

Fruit baskets, ping pong tables, "we're like a family," unlimited PTO (paid time off) that nobody actually takes – none of this tells you anything about the job. It's marketing.

Recruitment Agencies

Many job ads come from recruitment agencies, not the actual company.

Most mention this, but some don't. If the company name sounds generic ("Talent Solutions," "Tech Recruitment Partners") or there's no product or service mentioned, it's probably a recruiter filling a role for an undisclosed client.

This isn't necessarily bad – sometimes great roles come through recruiters. Just know what you're dealing with. The recruiter is a middleman with their own incentives. Respectively - don't do deep research into the recruitment company itself - it's meaningless.

Additional Research

Check for Other Roles at the Same Company

Sometimes companies hire multiple data roles at once across different teams. This is an opportunity.

Find them all. Compare the job descriptions.

- Are they consistent? (Mature org, clear thinking, probably good to work for)
- Are they wildly different? (Siloed teams, no central data strategy)
- Is there overlap in specific responsibilities? (Definitely ask about it)

This takes five minutes and tells you a lot about the company's data maturity. Also, the tech requirements are usually different, so you can figure out with pretty high certainty the specific tools that they use.

Red Flags Summary

Quick reference for things that should give you pause. Also check out the annexes below for specific examples:

- Job description assembled from multiple unconnected sources
- Vague responsibilities with no concrete examples
- Perks section longer than responsibilities section
- No clear reporting structure
- Role has been open for 6+ months (check LinkedIn)
- Previous person left after less than a year
- Emphasis on equity at early-stage startup
- "Fast-paced environment" (often means "chaotic and under-resourced")
- Requiring 100% in-office for a role that clearly could be done remotely

Summary

Reading job ads is a skill that improves with practice. Here's a simple process:

1. First pass: Skim for type. Is this capacity or capability? Which of the eight types does it resemble?

2. Second pass: Read the requirements. What's boilerplate vs. signal? Can you close the gaps quickly?
3. Third pass: Check the meta. Location, compensation signals, who posted it. Looks OK? You might want to save it for next steps at <https://myapplicationtracker.app/>
4. Research: Look at other roles at the company. Check LinkedIn for people in similar positions. How long do they stay?
5. Decide: Does this match what you're looking for? If yes, prepare questions for the interview that will verify your understanding.

Most people apply to too many jobs without thinking. Then they wonder why they end up in roles that don't fit. Or, more likely, don't end up finding a job at all. A little analysis up front saves a lot of pain later.

Your career is your choice. Have a system.

Annex: Real-World Example – Type Analysis, 1

The following is a real job ad for a Data Engineer position at a fintech company, with company name removed. This example demonstrates how to apply the framework from this guide to identify the job type.

Your Mission

[Company] builds technology that helps financial institutions detect and prevent financial crime. The systems you support are used every day to identify suspicious behaviour, stop bad actors, and help regulated companies meet strict compliance expectations.

As product usage has grown, so has the volume and importance of our data. Our teams rely on this data for analytics, product decisions, customer insights, detection quality, and operational workflows. Today, the product generates far more data than we can comfortably analyse or use.

Your mission is to build the data foundations that let the company use this information effectively. You will help internal analytics move quickly, support product teams with reliable insights, and unlock new data-driven value for customers.

Your Role

As a Data Engineer, you will:

- Design, build, and maintain data pipelines (today we are using S3, Glue, and Redshift).
- Improve the reliability, speed, and quality of data used by analytics, product, and operational teams.
- Work with existing tooling to deliver value quickly, while identifying longer-term improvements to the data platform.
- Collaborate with Analytics and Data Scientist to understand data needs and remove bottlenecks.
- Partner with Product and Engineering to ensure the data platform supports product insights, reporting, and new capabilities.
- Introduce technologies, patterns, and approaches when they provide clear value and support future scale.
- Over time, and if the impact is clear, this role may expand into shaping or leading a dedicated data engineering function.”

What does success look like after 6 months?

- Improve the internal analytics workflow by making key datasets easier and faster to access.
- Deliver the first reliable set of pipelines and models that reduce manual data preparation and improve consistency.
- Establish clearer structure, naming, and quality standards for core datasets used across the organisation.
- Identify the major bottlenecks in the current data platform and propose a realistic roadmap for addressing them.
- Produce early foundations of a scalable data platform (e.g. first redesigns, improved data models, clearer ingestion patterns).

The Analysis

This job ad is **Type D: Capability - New Function in an Existing Team.**

Why This is Type D

First, it's clearly a capability hire, not capacity. The mission statement says: "Your mission is to build the data foundations that let the company use this information effectively." They're not adding another person to do existing work – they need someone to build something new.

Second, an existing team is adding a genuinely new capability. The company exists, they have a product, they have analytics/product/engineering teams to collaborate with. But they explicitly state: "Today, the product generates far more data than we can comfortably analyse or use." They need someone to build the data platform that doesn't really exist yet.

Third, they provide specific business outcomes, but freedom in execution. They are clear on what success looks like (6-month goals are well-defined), but the "how" is up to you. They mention "future evolution expected" in orchestration, and tools like dbt, Airflow, or Dagster as "future potential."

Fourth, there's a growth signal: "Over time, and if the impact is clear, this role may expand into shaping or leading a dedicated data engineering function." This is classic Type D language – you're building something that could become a team.

Why it's not other types

Not Type A or B (Capacity): This isn't backfill or adding hands to existing processes. There's no mention of joining an existing data engineering team or following established patterns.

Not Type C (Capability Disguised as Capacity): The ad is very clear and coherent, not assembled by confused HR. The responsibilities make sense together and clearly point toward building something new.

Not Type E (Greenfield): You're not completely alone. There are existing teams and some infrastructure (S3, Glue, Redshift already in use). You'll be collaborating with Analytics, Data Scientists, Product, and Engineering.

Not Type F (Chaos Backfill): The tasks are coherent and make sense together. This doesn't read like it was assembled by asking four different people what the previous person did.

Good Fit For

Self-starters who can work without detailed instructions. People who enjoy having a clear business goal but freedom in execution. The guide says "Ideal for freelancers looking for a permanent position." You should be comfortable defining what the function looks like at this company.

Annex: Real-World Example – Type Analysis, 2

The following is a real job ad for a Senior Data Analyst position at a B2C/iGaming company, with company name removed. This example demonstrates a clean type and also illustrates domain-specific requirements.

Requirements:

- At least 3 years of experience in product/ data analytics;
- Strong proficiency in SQL (including advanced queries and window functions);
- Hands-on experience with BI tools (e.g., Tableau, Power BI);
- Understanding of relational database architecture (ClickHouse, PostgreSQL, BigQuery, etc.);
- In-depth knowledge of statistics and hypothesis testing methods;
- Skills in hypothesis validation and data-driven analysis.
- Will be plus:
- Experience with Data Science tools for automating A/B test analysis is a plus.
- Experience in iGaming or B2C related industry (preferably whale-driven).
- Responsibilities:
- Researching for insights, opportunities and growth points aligned with department KPIs;
- Micro-segmentation, offer personalization;
- Regular hypothesis backlog validation and prioritization;
- Customer behavior and product usage analytics (including product, website and mobile app data);
- KPI monitoring and alerting, finding 'why?';
- Developing and maintaining BI dashboards;
- Adhoc analytical requests (segmentation, data export, metric snapshots);
- Brainstorming and collaboration with different business streams;
- Data accuracy validation;
- Recurring analysis tasks automation to reduce manual overhead.

The Analysis

This job ad is **Type A: Capacity in a Mature Team**.

Why This is Type A

First, it is a clear **capacity hire**. The team knows exactly what they need: a safe pair of hands to step into an existing flow. The responsibilities - KPI monitoring, dashboard maintenance, ad-hoc requests - are the bread and butter of a running analytics function. They aren't asking you to invent the role; they are asking you to fill it.

Second, the **requirements are coherent and standard**. Advanced SQL, BI tools, and statistics are the classic toolkit for this role. There is no "wishlist" confusion here (e.g., asking an analyst to also manage Kubernetes clusters), which signals a mature management layer that understands the distinction between data roles.

Third, the focus is on **depth and optimization**. Phrases like "regular hypothesis backlog validation" and "recurring analysis tasks automation" suggest a team that has moved past the initial "chaos" phase. They have processes in place (a backlog, KPIs, specific segments), and they need someone to execute and refine them, not build them from scratch.

Why it's not other types

- **Not Type B (New/Small Team):** The mention of "department KPIs," "different business streams," and specialized tasks like "micro-segmentation" implies a larger, more structured organization, not a scrappy startup team of two.
- **Not Type C (Capability Disguised):** The link between the skills asked (SQL, Stats) and the work required (Hypothesis testing, Segmentation) is perfect. There is no disconnect here; they know what they are buying.
- **Not Type F (Chaos Backfill):** The responsibilities list is focused entirely on analytics. It doesn't veer into data engineering, IT support, or office management.

Good Fit For

People who value **stability and clarity**. If you want to come in on day one, be given access to a database that works, and start answering interesting questions, this is for you. It is ideal for someone who prefers **depth** (getting really good at iGaming analytics) over **breadth** (wearing 10 different hats).

Annex: Real-World Example – Type Analysis, 3

The following is a real job ad for a large healthcare corporation in the United States. It's a lot of text, but it's an exact copy-paste, just to demonstrate how some ads can look. Feel free to skip to the analysis below and then refer back to the ad itself.

Responsibilities

The Data Engineer gains access to data across the organization and provides ongoing analysis of the data by monitoring, profiling and analyzing databases. Requires a mix of functional, data and technical skills. The right candidate must be able to understand business requirements, translate them into information needs and implement those requirements using data available. The hire will be responsible for expanding and optimizing data architecture, as well as optimizing data flow and collection for cross functional teams. The ideal candidate is an experienced data pipeline builder and data wrangler who enjoys optimizing data systems. The Data Engineer will support our software developers, database architects, and data analysts and data scientists on data initiatives and will ensure optimal data delivery architecture is consistent throughout ongoing projects.

Core Job Duties Include

Assemble large, complex data sets that meet functional / non-functional business requirements.

Strong knowledge of SQL required. Ability to identify sets and subsets of information across multiple joins or unions of tables is preferred in addition to writing and troubleshooting SQL queries for data mining

Perform complex data analysis and investigation for customer requests to explain results and to make appropriate recommendations.

Strong understanding of data modeling concepts

Problem solver with the initiative to think critically to identify improvement opportunities (error detection, error correction, root cause analysis)

Understand ETL that will aid in verification and testing of data

Build processes supporting data transformation, data structures, metadata, dependency and workload management.

A successful history of manipulating, processing and extracting value from large disconnected datasets.

Analyze business objectives and develop data solutions to meet customer needs.

Demonstrated ability to effectively participate in multiple, concurrent projects

Improve and customize current data solutions to meet business functional and non-functional requirements.

Research new and existing data sources in order to contribute to new development, improve data management processes, and make recommendations for data quality initiatives.

Perform periodic data quality reviews for internal and external data.
Ensure timely resolution of queries and data issues.
Look for new ways to find and collect data by researching potential new sources of information.
Work with data and analytics experts to strive for greater functionality in our data systems.

This is a full time, remote position within the United States.

During the selection process you may participate in an OnDemand (pre-recorded) interview that you can complete at your convenience. During the OnDemand interview, a question will appear on your screen, and you will have time to consider each question before responding. You will have the opportunity to re-record your answer to each question - [company] will only see the final recording. The complete interview will be reviewed by a [company] staff member and you will be notified of next steps.

A Codility Test may be required as part of the candidate selection process.

Qualifications

Bachelor's degree in Computer Science, Engineering, or related field from an accredited University or College; OR an Associate's degree in Computer Science, Engineering, or related field from an accredited University or College with 2 years of experience. Demonstrated ability to analyze and profile data as a means to address various business problems through leveraging advanced data modeling, source system databases, or data mining techniques, is required. May provide consultative services to departments/divisions and committees. Demonstrated application of several problem-solving methodologies, planning techniques, continuous improvement methods, and analytical tools and methodologies (e.g. data analysis, data profiling, modeling, etc.) required. Incumbent must have ability to manage a varied workload of projects with multiple priorities and stay current on healthcare trends and enterprise changes. Interpersonal skills and time management skills are required. Requires strong analytical skills and the ability to identify and recommend solutions, advanced computer application skills and a commitment to customer service. Experience with data analysis, quality, and profiling; including data exploration tools including but not limited to Rapid SQL, AQT, Information Analyzer, and Informatics.

Experience with Python, Oracle OCI/FDI/ODI, Automation/AI Experiences, Google Cloud and/or Microsoft Azure data engineering preferred.

The Analysis

This job ad is **Type F: The Impossible Job, Confused Enterprise Edition**.

Why This is Type F

At first glance, this looks like a standard corporate job. But look closer at the "Frankenstein" nature of the responsibilities.

- **The Role Mashup:** They want a Data Engineer (build pipelines), a Data Analyst (explain results to customers), a Database Administrator (monitor/profile databases), and Support Staff (ensure timely resolution of queries). These are four different jobs.
- **The Copy-Paste Effect:** "Profiling" and "Data Quality" appear three separate times in different phrasings. "Business requirements" appears three times. This signals that the job description wasn't written by one person with a plan; it was assembled by HR pasting together wish-lists from different stakeholders who didn't talk to each other.
- **The "One-Man Band" Legacy:** This reads exactly like a backfill for someone who had been there for 5 years and "just handled everything." Now that they are gone, the company is trying to hire one person to replace a unicorn.
- **The Qualifications block functions as a literacy test:** if you can actually parse this wall of text, you might just have the patience required to survive the job itself.

Why it feels different from other Type Fs

Unlike a startup Type F (which is high-energy chaos), this is **bureaucratic chaos**. The presence of "OnDemand interviews," strict degree requirements, and specific tool lists (Rapid SQL, AQT) suggests a rigid environment. You won't be "running and gunning"; you will likely be stuck in the middle of a turf war between IT and Business, trying to fulfill contradictory requirements.

Good Fit For

The **"diplomat engineer."** If you are technically decent but world-class at managing stakeholders and saying "no" politely, you can survive here. The actual technical bar is likely low (maintenance mode), but the political bar is high. If you can navigate the bureaucracy, it's a secure paycheck. If you expect clear requirements and focused work, you will burn out in six months.

Annex: Real-World Example – Type Analysis, 4

This is a real job ad and I stopped reading at the line in bold below - “develop an understanding of our current data landscape”. But just for the purposes of analysing it I decided to continue.

The Data Engineer II will play a vital role in our Data transformation journey - helping move the organization toward a future state where decision-makers have the right data, in real time, when they need it. This role is responsible for driving the execution of key data projects, including data management, data conversions, ongoing ingestion pipelines, and data platform engineering and operations. The ideal candidate will work across systems and teams to ensure data flows seamlessly, reliably, and securely throughout the enterprise, supporting both strategic initiatives and day-to-day operations.

Responsibilities

- **Develop an understanding of our current data landscape**
- Administer and optimize Snowflake and other data platforms to ensure performance and alignment with business goals including role/permission management performance tuning, cost optimization, and capacity planning.
- Manage operations of the Data Warehouse, Middleware/Azure Enterprise Service Bus (ESB), Cognos, and related analytics tools.
- Provide secondary support for Data Integration & Platform Engineering responsibilities, including assisting with system maintenance, data pipeline operations, and platform enhancements.
- Stay informed on emerging data technologies and trends, sharing insights to help identify opportunities for innovation and improved data use.
- Help implement scalable and efficient data solutions that improve business insights and operational performance.
- Collaborate with IT and business stakeholders to support data access, usability, and consistency across departments.
- Participate in the integration of new tools and systems, ensuring seamless functionality and data accuracy.
- Provide advanced support by resolving escalated issues from Data Engineer I, ensuring timely and effective resolution of complex technical challenges.
- Contribute to a collaborative, high-performing team culture by supporting peers and sharing knowledge.
- Enforce data governance policies, ensuring data is secure, consistent, and compliant with regulations.
- Support efforts to meet financial regulations and cybersecurity standards in all data-related systems.
- Administrator and/or deliver configuration requests in our Azure ESB environment to Banking Core System through APIs.
- Design and develop pipelines for additional data vendors

- Recommend, implement and maintain best practices for the environment
- Ability to illustrate and maintain the enterprise's data systems architecture
- Complies with all applicable banking laws and regulations, including, but not limited to the Bank Secrecy Act, USA Patriot Act, and related anti-money laundering statutes, and federal consumer protection legislation and regulations. Builds working knowledge of all applicable laws and regulations.

Skills Required

- 3+ years managing data / ETL development
- 3+ years hands on in analyzing, designing which involve Data pre-processing , Data Extraction , Data Ingestion , Data Quality ,Data Normalization/Anonymizing & Data Loading
- 5 + years data analysis, Ad hoc reports, and visualization; Cognos, Tableau
- 5+ years data migration and validation projects
- 5+ years of Snowflake Management
- 5+ years implementation and support of on-prem and cloud DBs
- 5+ years' experience with system process mapping
- Prior work within the financial industry is preferred.

Education & Work Experience

- Degree- Any
- Work Experience- 4-8 years of experience

The Analysis

This job ad is **Type F: The Impossible Job, welcome to unicorn migration-land.**

Why This is Type F

The defining characteristic of this ad is the first line of the requirements. “Develop an understanding of our current data landscape”. This is the nicest way to say “we have absolutely no idea what’s going on”.

- **The Title:** "Data Engineer II." In most compensation bands, this is a mid-level role.
- **The Requirements:** They are asking for **5+ years** of experience in *everything*.
 - 5+ years in Snowflake (Cloud Warehousing).
 - 5+ years in Cognos (Legacy BI Reporting).
 - 5+ years in Data Migration.
 - 5+ years in On-prem and Cloud DBs.
 - Plus "Azure Enterprise Service Bus" (Backend Application Integration).
 - Plus work experience of 4-8 years. It's “pretty hard” to have 5+ years of experience in multiple separate technologies with just 4 years of work experience.

The "Frankenstein" Tech Stack

This isn't just a Data Engineering role; it is three jobs in a trench coat:

1. **The Modern Cloud Architect:** "Administer and optimize Snowflake... capacity planning."
2. **The Legacy IT Support:** "Manage operations of Cognos... Administrator Azure ESB."
3. **The Compliance Officer:** "Complies with Bank Secrecy Act, USA Patriot Act."

The "Admin" Trap

Notice the word "Administer" appearing twice. A Data *Engineer* builds pipelines. A Data *Administrator* manages permissions, costs, and uptime. This role is explicitly asking you to be the **Systems Administrator** for their entire stack. Someone started to modernise the stack, moving from on-prem to cloud, and then left.

Good Fit For

No one. Unless you have experience with this exact stack - then you're automatically hired and can probably get compensated very well. Even then, there are definitely several complex migrations halfway done in a compliance-heavy environment and then abandoned - this is not just keeping the lights on.