GUIDE 02

Builder or Steward? How Senior Data Leaders Are Read

Why strong careers are read the wrong way, and what that reading costs senior data and AI professionals in the GCC.

02

Kaleem Chaudry
For senior data and AI professionals in the GCC

Contents

INTRODUCTION	The same career, two readings
CHAPTER 1	What GCC buyers are paying for now
CHAPTER 2	How a steward is read
CHAPTER 3	The altitude problem
CHAPTER 4	When range reads as dilution
CHAPTER 5	Advisers and operators
CHAPTER 6	The cost of staying unread
CHAPTER 7	What this means for you
ABOUT	About the author

The cases in this guide are composites built from patterns observed across many careers. None describes a real individual.

INTRODUCTION

The same career, two readings

Picture two senior data leaders with almost identical records. Twenty years each. Similar institutions, similar team sizes, similar programmes delivered. Both are strong. Both are well regarded by the people who have worked with them.

One is approached for a role leading a new data and AI agenda. The other is approached, when approached at all, for roles running what already exists.

The difference is not in what they did. It is in how what they did is read. One is read as someone who builds. The other is read as someone who looks after things. In a market that is buying the first far more than the second, that difference shapes everything that follows: which roles come their way, at what level, and how seriously they are taken in the room.

This guide is about that reading. It covers what senior buyers in the GCC are paying for now, how experienced professionals come to be read as stewards rather than builders, and three related patterns that cause strong careers to be misread: narratives pitched at the wrong level, range that reads as dilution, and advisory experience that does not translate into operator credibility.

The emphasis is on consequences: what each pattern costs, and why. It does not set out how to reposition a career, because that depends entirely on the individual record. What it will do is help you see which of these readings may be attached to you.

As in the other guides in this series, each chapter opens with a short case. These are composites built from patterns seen across many careers. None describes a real individual.

CHAPTER 1

What GCC buyers are paying for now

> **COMPOSITE CASE**
>
> A senior leader at a regional financial institution was asked by her chief executive to brief the board on the organisation's AI ambitions. She had run its data governance function for eight years, and run it well. Policies were in place, controls were sound, audits were clean.
>
> The board listened politely. Afterwards, the chief executive told her the board had been looking for someone to tell them what the organisation could build and what it would be worth. They had heard, instead, about what was being protected. Six months later, a new role was created above her to lead the AI agenda. She was not considered for it.

To understand how senior data leaders are read, it helps to start with what the buyer is trying to buy. In the GCC, that has changed significantly over the past decade.

From control to creation

For much of the 2010s, senior data roles in this region were shaped by regulation and risk. Organisations needed governance frameworks, data quality controls, regulatory reporting and audit readiness. The leaders who delivered those things were valuable, and many built strong careers doing so.

That work still matters. But it is no longer what most senior buyers are excited to pay for. National strategies across the region now place data and AI at the centre of economic and government transformation. Boards and chief executives are under pressure to show what data and AI will create: new services, faster decisions, lower costs, new revenue, better outcomes for citizens and customers.

The buyer's question has shifted from "will this person keep us safe?" to "what will this person build, and what will it be worth?"

The questions buyers now ask

The shift shows up most clearly in the questions senior decision-makers ask. What did you create? What was it worth? How quickly did it reach the people who use it? What could the organisation do afterwards that it could not do before?

Candidates who answer those questions with process, describing frameworks, committees, policies and controls, are giving accurate answers to a different question. The decision-maker hears competence, and files the candidate under process. Candidates who answer with outcomes are filed under value. Both may have done the same work.

What this means for senior candidates

A senior hire is now usually assessed against a transformation agenda, even when the role includes governance responsibilities. Decision-makers look for evidence that a candidate has created value, led change and delivered outcomes the business can feel.

Professionals whose experience is genuinely in that space can still be read the wrong way if they describe it in the language of the earlier era. Equally, professionals whose careers were built in governance and control often did far more value creation than their materials suggest. They simply never needed to describe it that way.

The consequence

When a candidate's narrative speaks to what the market used to buy rather than what it buys now, the result is rarely outright rejection. It is quieter than that. The candidate is considered for a narrower set of roles, at a lower level of ambition, and the transformation mandates go to someone whose story matches the agenda.

CHAPTER 2

How a steward is read

COMPOSITE CASE

A Head of Data Management at a large government-related entity applied for a Chief Data Officer role at a fast-growing private group. Over fifteen years he had built a data platform from almost nothing, consolidated systems across several business units, and given the organisation its first reliable view of its customers.

His CV described his role as "responsible for data governance, data quality and data management standards across the enterprise". The private group's chief executive read it as a compliance profile. The feedback was that he was "very solid", and that the group needed "more of a builder". He had been the builder. Nothing on the page said so.

The steward pattern is the most common positioning problem among senior data professionals in this market, and it is rarely about the quality of the work. It is about how the work is described, and therefore how it is read.

What the steward reading sounds like

A steward narrative is built around responsibility. It describes what a person was in charge of, what they oversaw, what frameworks they maintained and what standards they enforced. It is accurate. It is often modest. And it consistently tells the reader that this person looks after things.

A builder narrative is built around change. It describes what was created, what improved and what the organisation could do afterwards that it could not do before. It is often describing exactly the same work.

AN ILLUSTRATION

READ AS A STEWARD	*"Responsible for enterprise data governance, data quality standards and policy compliance across all business units."*
READ AS A BUILDER	*"Built the data foundation that let the organisation launch its first analytics and AI services, and gave leadership a single trusted view of its customers."*

Illustrative only. The same career can produce either line. Which builder framing is true, and supportable, depends entirely on an individual's real outcomes.

What a decision-maker infers

A senior reader does not consciously label a candidate as a steward. They form an impression, quickly, and the impression carries assumptions with it.

They assume the person is reliable, careful and thorough. They also assume, often without realising it, that the person is better suited to maintaining an existing function than to building a new one. They wonder whether the person will slow things down. They see a cost centre rather than a source of value. And they quietly place the candidate in a category that is respected but rarely chosen to lead change.

The consequences

The steward reading costs senior professionals in specific, recurring ways.

The wrong roles arrive. Approaches, when they come, are for governance, controls and operations roles, often at the same level as the person's current position or below it. Transformation and AI leadership roles go elsewhere.

The ceiling lowers. Even when a steward-read candidate is hired into a senior role, the scope offered tends to be narrower and the mandate more defensive. The person is hired to keep something safe, not to build something new.

Internal standing erodes. The pattern does not only operate in the external market. Inside organisations, leaders read as stewards are often passed over when new data or AI agendas are created, and a new role appears above them, as in the case at the start of Chapter 1.

Confidence follows. Many capable leaders conclude, after a period of being overlooked, that the market has moved past them. In most cases it has moved past their description, not their ability.

CHAPTER 3

The altitude problem

> **COMPOSITE CASE: TWO PROFESSIONALS**
>
> The first was a Senior Manager targeting Head of Function roles. In interviews, he spoke in detail about milestones, tooling choices and delivery mechanics. Every answer was competent. Every answer confirmed that he was a strong Senior Manager.
>
> The second was a Head of Function targeting C-level roles. Her narrative was full of vision, strategy and enterprise transformation. When the chief executive asked what she had personally owned, in budget and decision rights, the answer was thinner than the narrative. She was read as reaching beyond her evidence.

Every senior band buys something different. A narrative pitched at the wrong band is misread even when the underlying record is sound. This is the altitude problem, and it fails in both directions.

What each level is buying

Broadly, the market reads senior data and AI roles at three altitudes.

Senior Manager roles buy delivery excellence and team leadership: someone who gets complex work done well through others.

Head of Function roles buy mandate ownership: someone who holds a budget, sets a direction for a function, and influences peers and leadership.

C-level roles buy enterprise value and board fluency: someone who connects data and AI to the organisation's strategy, and can represent it credibly to a board, a regulator or a government stakeholder.

A narrative signals altitude through what it chooses to emphasise. Detail about how things were done signals one level. Ownership of outcomes and decisions signals another. Strategic framing signals a third.

Titles do not travel cleanly

In the GCC, altitude is complicated by titles. A Director at a global bank, a Director at a government entity and a Head of Function at a regional group can hold very different scope. Some organisations use senior titles generously; others are conservative. Professionals moving between entity types often find that their title signals a different level to the new reader than it did at home.

Experienced readers know this, so they look past the title to what was actually owned. Candidates who rely on the title to carry their level are misread in both directions: overestimated by some readers, underestimated by others, and rarely placed accurately.

Pitched too low

The more common version is a narrative pitched at the level a person currently holds, not the level they want. It is a natural result of describing a job accurately. But the reader is not asking whether you do your current job well. They are asking whether you are already operating at the next level.

The consequence is that a strong candidate is read as "not quite ready". The feedback is usually vague, and the candidate often responds by adding more detail, which reinforces the reading.

Pitched too high

The less common but more damaging version is a narrative that claims a level the evidence cannot support. Vision without ownership, strategy without budget, enterprise language without enterprise outcomes. A sophisticated reader notices

quickly, and the consequence is worse than being read as not ready. The candidate is read as overstating, and that affects how everything else they say is received.

The consequence

The altitude problem quietly decides which band a candidate competes in. Pitched too low, they compete for roles they have already outgrown. Pitched too high, they lose credibility with the very people they need to convince. In both cases, the record itself was never the issue.

CHAPTER 4

When range reads as dilution

> **COMPOSITE CASE**
>
> A data and analytics leader had a career most people would envy: data governance, then analytics, then a data platform build, then a spell leading a digital programme, then risk analytics. Twenty-two years, five distinct strengths, every one of them genuine.
>
> When a search firm asked one of his former colleagues to recommend the best data platform leader she knew, she named someone else. When asked about him afterwards, she said he was "great at everything, but I'm not sure what I'd call him". He was never put forward for a role he would have been ideal for.

Long, varied careers are common among senior data professionals, and the range is real. The difficulty is that the market rarely buys range at senior level. It buys a clear strength that solves a specific problem.

Why range reads as dilution

When a reader cannot identify one thing a candidate is known for, they fill the gap themselves. The most common conclusion is "generalist", and in a market buying transformation, a generalist reads as a safe pair of hands rather than the person to lead a specific agenda.

Several themes competing in the same narrative also weaken each other. Each one gets less space, less evidence and less emphasis. The reader comes away with a broad impression and no sharp one.

The connection to how you are found

This pattern links directly to the referral route described in Guide 01. When a search consultant asks for the best person in a specific area, people recommend whoever comes to mind first for that area. A professional known for one clear thing is easy to name. A professional known for many things is often not named at all, even by people who rate them highly.

A market-specific twist

In the GCC, this pattern is sharpened by the way many organisations hire. Senior roles are often defined around a specific, current mandate, and the market tends to hire for the specialism that matches it. Accumulated range, however valuable it may be in practice, is easily discounted in that process.

Range is not the problem

None of this means that a varied career is a weakness. At the most senior levels, breadth can be exactly what an organisation needs. The difference is that, at those levels, breadth is presented as one coherent story about enterprise value rather than as a list of separate strengths. Range becomes a problem only when the reader is left to assemble it themselves.

The consequence

The cost of dilution is invisibility at the moment that matters. Range does not lead to rejection. It leads to not being named, not being shortlisted and not being remembered for the role that fits. Highly capable people are overlooked precisely because they are capable of so much.

CHAPTER 5

Advisers and operators

COMPOSITE CASE

A director from a large consulting firm had spent twelve years advising GCC clients on data strategy and digital transformation. He knew the region, knew the language and had worked with several of the entities he now wanted to join. He expected to step into a senior in-house transformation role without difficulty.

Eight months later, he had been interviewed several times and offered nothing. The feedback, when it came, was consistent: impressive, articulate, clearly knowledgeable. Not clearly someone who had owned the outcome. One hiring executive put it bluntly: he had designed many transformations, but who had he been accountable to when one went wrong?

Consulting backgrounds are common among senior data and AI candidates in the GCC, and many consultants assume that their experience will transfer directly into in-house leadership. In my conversations with professionals making this move, that assumption comes up again and again. In practice, it is rarely that simple.

What a consulting title does and does not prove

A senior consulting title proves real things: exposure to many organisations, strong analytical and communication skills, familiarity with transformation methods, and the ability to work with senior stakeholders.

What it does not prove, on its own, is ownership. An in-house executive is accountable for outcomes over years. They hold a budget, carry a team, live with the consequences of decisions and answer for results when the advisers have left. Hiring executives know the difference, because many of them have been on the receiving end of consulting work.

Why the gap is wider in the GCC

Many senior hiring executives in this region have worked extensively with consulting firms on national and organisational transformation programmes. They are experienced buyers of advice. That experience makes them particularly alert to the difference between someone who recommended a transformation and someone who delivered one.

There is also a question of what translates. Experience that is valued in a consulting context, such as breadth of clients, frameworks and engagement leadership, does not automatically carry the same value in an in-house role in the GCC, where the buyer wants evidence of sustained delivery, accountability and results inside an organisation like theirs.

The language of advice

Consulting careers also carry a characteristic vocabulary. Advised, supported, recommended, designed, led the engagement. Each word is accurate, and each one tells the reader that someone else held the outcome. Across a whole CV or a whole conversation, the cumulative effect is a clear signal: this person has been close to transformation, but on the outside of it.

It is not a question of disguising a consulting background. Experienced readers would see through that immediately, and the background has real value. It is a question of whether what the person genuinely owned is visible at all.

What a decision-maker infers

When a consulting narrative leads with clients, engagements and recommendations, a hiring executive tends to draw several conclusions. This person understands the problem. They may not have lived with the solution. They may struggle with the slower pace and the politics of an in-house role. They may leave for the next interesting engagement. None of these may be true. All of them are common inferences.

The consequence

The consequence is a pattern many former consultants recognise: strong interest, positive conversations, and offers that do not arrive, or arrive at a lower level than expected. The title opened the door. It did not carry the candidate through it. Where genuine operator experience exists, such as periods spent in-house, accountability for delivery, or responsibility for a budget or team, it is often buried beneath the advisory story instead of leading it.

CHAPTER 6

The cost of staying unread

> **COMPOSITE CASE**
>
> A senior analytics leader recognised, in hindsight, that her positioning had been working against her for about five years. In that time she had been approached for roles at her existing level, been passed over twice for internal promotions into new AI roles, and turned down one external role she wanted because the scope offered was smaller than her current job.
>
> None of these events felt connected at the time. Each had its own explanation. Together, they formed a pattern that had cost her a significant part of her most productive career years.

The patterns in this guide rarely produce a single dramatic setback. Their cost is cumulative, which is exactly why they are so easy to overlook.

How the cost compounds

Roles. The roles that arrive are shaped by how a person is read. A steward reading attracts stewardship roles. An under-pitched narrative attracts roles at the current level. A diluted story attracts nothing in particular. Each role accepted on those terms reinforces the reading for the next move.

Level. Every move made at the wrong altitude makes the next step up harder to justify. The gap between where a person is and where their ability would place them widens quietly over time.

Time. Senior careers have a finite number of significant moves left in them. A positioning problem that persists for several years can consume one or two of those moves.

Confidence. Perhaps the most damaging cost is internal. Repeatedly being overlooked leads capable people to doubt their relevance. They begin to accept the market's reading of them as accurate, and to lower their own ambitions to match it.

Why it is noticed late

The cost tends to be recognised only in hindsight, for a simple reason: the market does not give feedback on positioning. Nobody writes to say that a candidate was read as a steward, pitched at the wrong altitude or remembered for nothing specific. Recruiters and hiring executives give polite, general reasons, or none at all. The professional is left to interpret silence, and most interpret it as something about themselves rather than something about how they are described.

What this cost is not

It is not a verdict on ability. In almost every case described in this guide, the underlying record was strong. What failed was the reading, and a reading can be changed in a way that a track record cannot.

That is the encouraging part of this guide. Of all the gaps senior professionals face, positioning and narrative is among the most common, and among the most addressable, because the raw material is usually already there.

CHAPTER 7

What this means for you

The patterns in this guide share one feature: they turn a strong record into a weaker reading. A builder described as a steward. A future Head of Function described as a strong Senior Manager. A specialist with real depth described as a generalist. A consultant with real operator moments described as an adviser.

Each reading carries consequences: narrower roles, lower levels, lost time and eroded confidence. None of them reflects what the person is actually capable of.

Do you recognise any of these patterns?

The questions below are not a diagnosis. They are a way of noticing which readings may be attached to you.

- [ ] Does your profile mainly describe what you were responsible for, rather than what changed because you were there?
- [ ] Are the roles that reach you mostly about running, governing or maintaining, rather than building?
- [ ] Has feedback described you as "very solid" or "a safe pair of hands" when you were hoping for something else?
- [ ] Does your narrative describe the job you have now, rather than the one you want next?
- [ ] If a former colleague were asked what you are best known for, would they give one clear answer?

- [ ] If you come from consulting, does your story lead with clients and engagements rather than outcomes you owned?
- [ ] Have new data or AI roles been created around you, or above you, without you being considered?

If several of these felt familiar, you are far from alone. These are among the most common patterns in senior data and AI careers in this region.

Where this fits

This guide covers Positioning & Narrative, one of the four pillars in the Senior Career Positioning Benchmark. The others are Market Visibility, Proof & Credibility, and Market Leverage & Direction. Narrative sits at the centre of them: it shapes how you are found, how your evidence is interpreted and what roles you are considered for.

Recognising the pattern is the starting point. Changing the reading is specific to each record: which outcomes are real, which framing is supportable, and which level the evidence will carry. If you want help applying this to your own record, the Benchmark is where that starts. If you have already taken it, re-read your Positioning & Narrative result with this guide in mind.

ABOUT THE AUTHOR

Kaleem Chaudry

Kaleem has spent over 25 years across JP Morgan, UBS and ADIA, with 360-degree GCC recruitment experience on both the hiring and candidate side. He now works with senior data and AI professionals on how they are positioned, evidenced and found in the GCC market.

He is the creator of the Senior Career Positioning Benchmark, a structured self-assessment for senior data and AI leaders based in the GCC or targeting a move to it.

This guide provides general information only. The case studies are composites built from patterns observed across many careers; they do not describe real individuals. The illustration in Chapter 2 is generic and is not a template. Nothing in this guide creates recruitment representation, provides access to vacancies, or guarantees interviews, offers, salary outcomes or employment.

© Kaleem Chaudry Consulting Services