THE EXPENSIVE ILLUSION OF DATA MATURITY:

Why most Enterprise Data Investments are still not delivering

Over the past decade, enterprise leaders have poured extraordinary capital into the promise of becoming data-driven organisations. The narrative has shifted with each wave of technology.

Between 2015 and 2019, the rallying cry was consolidation: get all your structured data into a single warehouse, build a foundation for reporting and business intelligence.

From 2020 to 2023, the ambition expanded to data lakes, with organisations racing to capture both structured and unstructured data in cloud storage.

And **from 2023 onward,** the conversation has pivoted sharply towards artificial intelligence, with boards asking why their data science teams haven't yet delivered transformational outcomes.

At each stage, the vision was compelling.

And at each stage, the majority of organisations ended up in the same place: spending more, governing less, and struggling to translate platform investments into business value.

Data warehouses became reporting engines that never matured beyond operational dashboards.

Data lakes became data swamps, vast and ungoverned repositories where cost grew far faster than insight.

And the rush to hire data scientists exposed a painful truth: most enterprise data platforms simply weren't ready for advanced analytics, leaving talented teams working from static extracts and one-off models that never made it into production.

The numbers bear this out. In our 2025 Data Maturity Survey of 136 organisations, **only 15 per cent scored themselves above 80 per cent** for achieving real analytics scale. The striking finding wasn't that companies lacked ambition. Most scored strongly on vision and strategic intent. **The gap was in execution: the platforms, capabilities, and governance programmes needed to turn ambition into outcomes consistently fell short.**

In our 2025 Data Maturity Survey of 136 organisations:

only 15% scored themselves

above 80% for achieving real analytics scale

The three patterns that keep repeating

When we examine why organisations stall, the same three patterns appear with remarkable consistency across the industries, geographies, and technology landscapes we work in.

The first is technology-first thinking.

Organisations select a platform, whether Fabric, Databricks, Snowflake, or something else, before they have clearly defined the business outcomes they need to achieve. The implicit assumption is that if you build a powerful enough platform, value will follow. It rarely does.

Without a clear connection between the technology investment and specific, measurable business objectives, platforms become expensive shelfware.

Users don't understand what the platform offers or how to access it, executive sponsors lose patience, and the initiative quietly stalls.

The platform may be technically excellent, but technical excellence without business alignment is just cost.

The second pattern is treating data as IT infrastructure.

When data is owned entirely by technology teams, it gets managed like plumbing: pipes, storage, ingestion jobs, and schema design. What's missing is business context. Data engineers build pipelines without understanding the commercial decisions those pipelines are meant to support.

Business teams, in turn, see data as someone else's problem. Nobody owns data quality at the domain level. Nobody is accountable for whether the definitions in a report actually match how the business operates. The result is a growing gap between what the platform contains and what decision-makers actually trust and use.

The third pattern is the big-bang transformation.

An organisation commits to an 18-month, enterprise-wide data transformation. Scope is ambitious, timelines are optimistic, and the implicit promise is that everything will come together at the end. But the end keeps moving. Requirements evolve as the market shifts. Scope expands as stakeholders add new demands.

Budgets overrun. And critically, the business sees no value until the entire programme is supposedly complete, which means political capital and executive patience are exhausted long before delivery.

By the time the programme is quietly wound down or restructured, the organisation has spent millions and has little to show for it beyond a partially built platform and a deeply sceptical business.

Underneath all three patterns lies a single, fundamental issue: **data was treated as a byproduct of systems and processes, not as a product in its own right.** Until that mindset shifts, no amount of technology investment will close the gap.

What organisations that succeed **actually do differently**

The organisations we've seen achieve genuine enterprise-scale analytics didn't get there by accident, and they didn't get there by buying better technology. They got there by making a deliberate, sustained decision to build an enterprise-grade data platform.

And they understood from the outset that "enterprise-grade" means far more than choosing the right cloud vendor. They start by treating data as a product with real customers, not as exhaust from operational systems. This means defining clear ownership, intentional design and a lifecycle for each key dataset or analytics asset: who it serves, what decisions it enables, how relevant/accurate it is and how users discover/access it.

Instead of one-off reports and ad hoc extracts, successful organisations create curated, well-documented data products that are designed, versioned, tested and supported just like any other product in the business.

This product mindset forces teams to prioritise usability, reliability and trust while still naturally connecting governance, architecture and business value into a single unified way of working.

These organisations invest in four interconnected pillars.

The first is unified governance.

A comprehensive approach to cataloguing data, tracking lineage, enforcing quality standards, and ensuring compliance and security across every data asset. Governance isn't a project that happens after the platform is built; it's embedded from the beginning.

The second pillar is production operations.

Successful organisations treat their data platform with the same rigour as a production software environment, implementing CI/CD pipelines, automated monitoring and alerting, and continuous performance optimisation. Data pipelines are not scripts running on someone's laptop. They are engineered, tested, deployed, and monitored like any other critical business system.

The third pillar is cost management.

FinOps practices are embedded into the platform from day one, with continuous resource optimisation, consumption controls, and budget governance that ensures the investment is delivering more value than it costs.

The fourth pillar is scalable architecture.

Modular, built from reusable components, and designed with an AI-ready foundation that anticipates where the organisation needs to go, not just where it is today.

What's important to understand is that these organisations don't start with a massive enterprise rollout. They start small. They prove value in focused, time-boxed sprints, often as short as eight weeks, targeting specific business use cases with quantifiable outcomes.

But from the very first sprint, they are building with scale in mind. Every component is designed to be reused.

Every governance decision is made with the enterprise context in view. And once they have proven success, they deliberately invest in expanding the foundation rather than simply replicating point solutions across teams. This is the critical difference between organisations that have five disconnected dashboards and organisations that have a governed, trusted platform serving hundreds of users across every function.

The evidence from our client engagements reinforces this.

A major food manufacturer modernised its analytics by building a unified whilst enabling self-service analytics with strong governance, establishing the foundation for predictive and augmented analytics across its complex supply chain.

A diversified chemicals group moved from fragmented legacy systems to a governed lakehouse architecture, delivering trusted data products, automated pipelines, and formal data quality governance across its mining, agriculture, and chemicals divisions.

And a mega-scale urban development programme achieved 126 per cent faster access to business-ready data, a 90 per cent reduction in manual reconciliation effort, and onboarded over 500 users to a self-service analytics platform, all built on a single governed environment. In each case, success came not from the technology alone, but from the deliberate decision to build a governed, scalable, business-aligned platform.

How we help organisations **accelerate**

For leaders reading this and recognising their own organisation in these patterns, the path forward doesn't require another **18-month transformation programme.** It requires honest assessment followed by focused, deliberate action.

This is precisely the approach we take with our clients. We begin with a comprehensive diagnostic assessment, led by our specialist certified data architects, to establish a clear baseline of where your platform stands today and define a target-state architecture that reflects where the business needs to go.

This isn't a technology exercise. It's a strategic one, and we ensure it is shaped jointly by business and technology leadership. From there, we work across four acceleration tracks tailored to what your organisation needs most.

Our Enterprise Grade Roadmap provides the strategic programme to move from project-level solutions to a true enterprise platform, including business use case development and change management across the organisation.

Our FinOps practice creates transparency around what your data platform actually costs and whether the value it generates exceeds that cost, with proven optimisation approaches specific to data platforms.

Many organisations are shocked to discover how much they're spending on compute, storage, and licensing for data that nobody uses.

Our Governance Accelerator deploys frameworks such as Unity Catalog or Microsoft Purview in weeks using tested templates and use cases, embedding data governance into the platform itself rather than treating it as a separate workstream that runs in parallel and never quite connects.

Our Rapid Implementation track leverages Databricks and Microsoft Fabric accelerators to deliver an enterprise-ready lakehouse platform, fully configured to scale, with production-grade operations including CI/CD pipelines, DevOps, DataOps, and continuous audit built in from the start.

Every engagement is grounded in the same principle: we start with business-led use cases and quantified value, execute in focused sprints that deliver quick wins and quarterly business outcomes, and build governance and scalability into the foundation from day one.

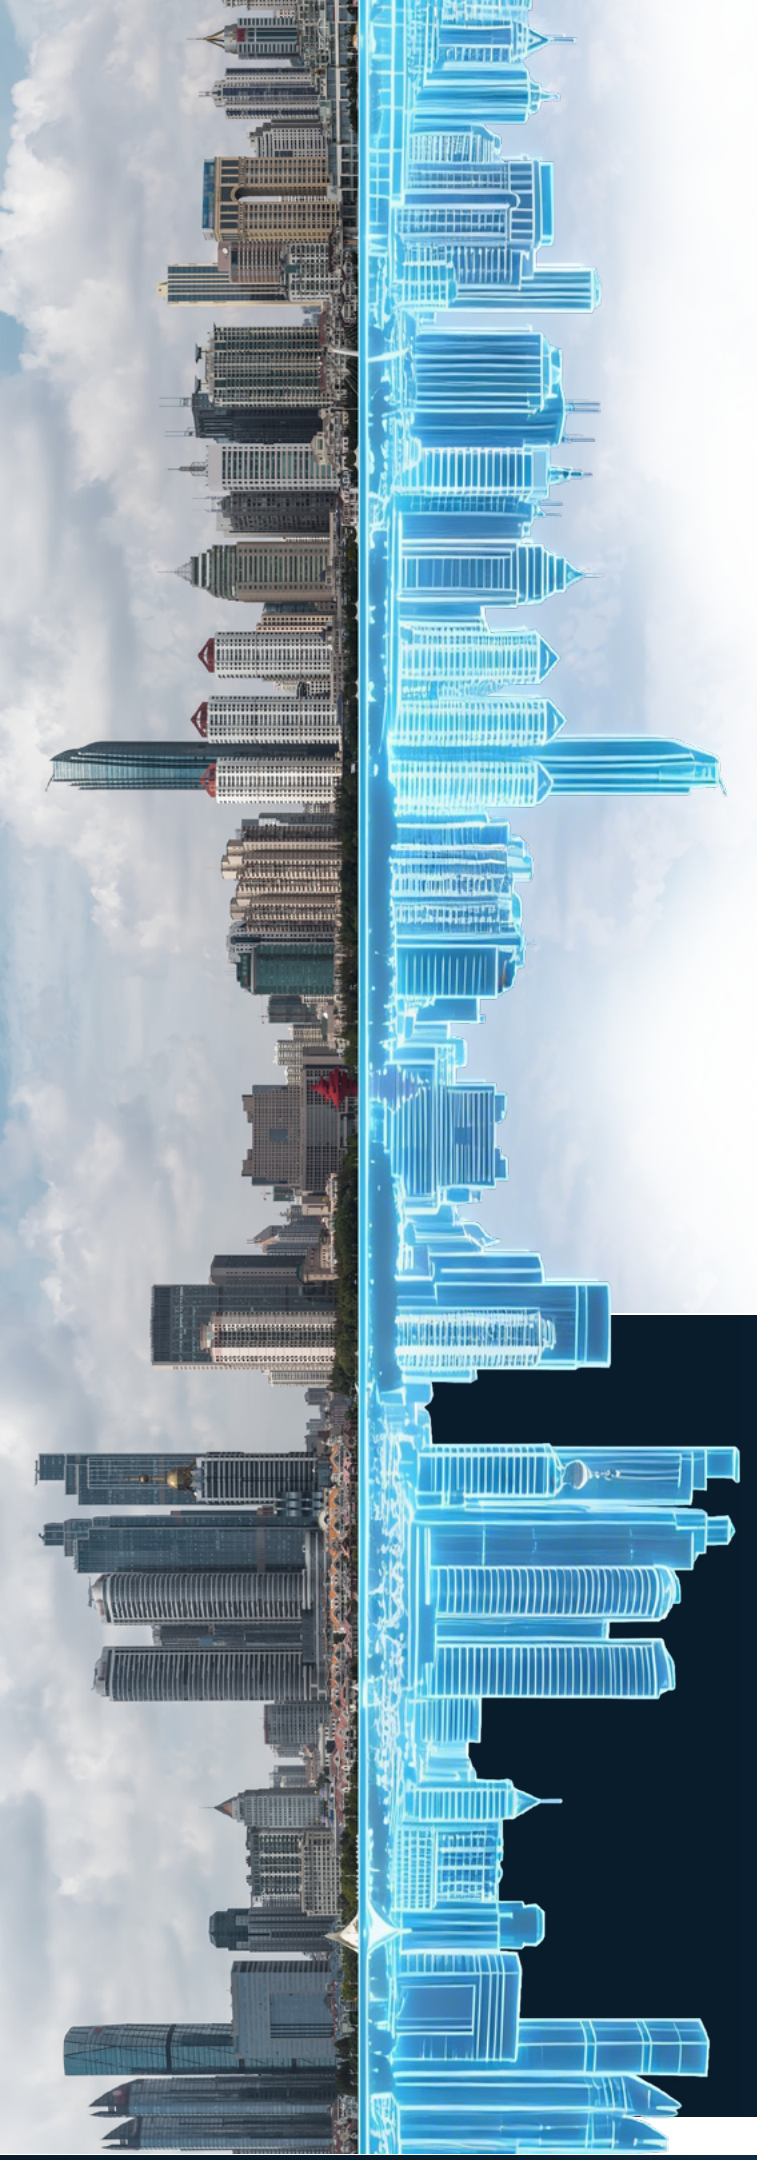

The urgency **IS REAL**

The window for getting data foundations right is narrowing. As AI capabilities accelerate, the gap between data-mature organisations and those still struggling with ungoverned, fragmented platforms will widen dramatically.

AI doesn't fix bad data.

IT AMPLIFIES IT.

Organisations that lack trusted, governed, well-architected data foundations won't just miss the AI opportunity; they'll fall further behind competitors who have already built the infrastructure to move quickly, experiment confidently, and scale what works.

The question is no longer whether to invest in your data platform. The question is whether you're investing in a way that will actually deliver.

DECISION INC.

Decision Inc. is an advisory-led cloud and AI partner that works with over 300 organisations annually to build and scale enterprise data platforms across financial services, retail, telecoms, manufacturing, and the public sector.

To learn more about how we can help accelerte your analytics journey, visit

www.decisioninc.com