

EXECUTION EDGE: **Data Foundations Before AI**

Start now. Start properly.



Executive Summary

AI is changing how businesses operate, compete, and create value. Its impact will deepen, not diminish. The winners will be those who move early and wisely with clear objectives, tight scope, strong processes and sound data foundations. The losers will rush to tools, deploy pilots at the edges and discover, again, that technology without foundations rarely delivers.

AI is different from past technology waves in its potential to reshape work and decision-making. But that potential is realised only when applied to well-designed processes and high-integrity data.

In digitally native environments, foundations often exist by default. In 10–20-year-old businesses with fragmented systems and messy processes, AI amplifies noise unless those foundations are deliberately built. Crucially, foundations need not be enterprise-wide. A scoped, domain-level foundation aligned to a specific objective is the fastest route to value and the best protection against expensive disappointment.

This paper sets out how to start: define the business objective, stabilise processes, establish data foundations for that domain and then introduce AI where it can demonstrably improve predictability, speed or cost.

Start now, but start properly.

The private equity lens is practical: value within a hold period, risk managed and story credible at exit. AI can support all three, if it improves core process outcomes with reliability.

In mid-cap service businesses, the path to that reliability begins with clarity and scope. "Don't eat the elephant" is not conservatism; it is strategy. Choose a thin slice with material impact, fix the process and the data that feed it, and use AI to accelerate what's already working.

Past technology waves such as ERP, CRM, big data and "digital transformation" promised step-changes. Many under-delivered not because the tools were weak, but because the inputs (process clarity, data quality, operating discipline) were not there. AI raises the stakes.

It is probabilistic, data-hungry and unforgiving of ambiguity. When applied to messy processes with inconsistent data, it produces confident wrong answers faster. When applied to well-defined work with trustworthy data, it compounds value.



What's different about AI and what isn't

What's different is the breadth of tasks AI can influence: classification, extraction, drafting, prediction, routing and decision support.

What isn't different is the dependency on process and data. AI systems don't magically "understand your business". They reflect whatever your processes and data encode.

Three properties matter:

- + **Sensitivity to data integrity:** AI performance degrades sharply with inconsistent definitions, missing fields, duplicate entities and stale records.
- + **Path dependence:** Early choices; definitions, schemas and integration patterns lock in costs or unlock speed for years.
- + **Probabilistic outputs:** You must design for verification, exception handling and human-in-the-loop where outcomes carry risk.

Start now, but start properly

Begin with a business outcome, not a technology exploration. Bound the scope to a domain where a measurable improvement matters: pricing adjustments, claims triage, invoice matching, service ticket routing, cash forecasting, demand planning.

For that domain:

- + Write down the **current process** clearly enough that two teams would execute it the same way.
- + Identify, then remove, the **variations** that create most rework.
- + Define the **target outcome** and how you'll measure it (accuracy, cycle time, cost per transaction, NPS impact, cash improvement).
- + Establish a **data contract** for the inputs and outputs: what fields, what types, what definitions, what timeliness.

Only then design the AI intervention.

The implication is straightforward: stabilise the process, standardise the data, then apply AI.



The data foundations: what “good” looks like

AI is only as good as the data it feeds on. If the data is messy, inconsistent, or incomplete, AI will magnify the problem rather than solve it. Strong data foundations are therefore not a “nice to have” they are the essential first step in making AI pay back.

So what does “good” look like?

Four things that matter the most:

1. Clarity of definitions

Everyone in the business should mean the same thing when they say “active customer,” “churn” or “gross margin.” Without this, reporting becomes inconsistent and AI models draw the wrong conclusions. A shared vocabulary and consistent definitions are the bedrock.

2. Quality and ownership

Data needs to be complete, accurate and up to date. That doesn’t mean perfect everywhere, but it does mean good enough where AI is being applied. Just as important, there must be clear accountability, someone responsible for making sure the data stays that way.

3. Accessibility and structure

Data that lives in isolated spreadsheets and legacy systems is hard to use. Bringing it together in a structured way so that teams can access the right version quickly is key.

This doesn’t have to mean building a huge enterprise platform. It simply means ensuring the data you need for the chosen process is organised, consistent and ready to use.

4. Trust and governance

People need confidence that the numbers are right and that sensitive information is handled appropriately. Light-touch governance, clear rules on who can access what, and how changes are made protects both trust and compliance.

Why this matters for AI: AI thrives on consistent, reliable data. With clean definitions, clear ownership, structured access and basic governance in place,

AI tools can produce insights and automation that are genuinely valuable. Without them, AI produces slick-looking outputs that quickly lose credibility.

Process first because process shapes the data

Messy processes produce messy data; messy data produces messy AI. Messy AI delivers unreliable results and potentially dangerous decisions. **Stabilising the process is not busywork it is the single biggest multiplier of AI impact.**

Focus on these three aspects:

1. **Decision points:** Where are decisions made? What are the business rules? What constitutes a “good” decision? Encode those criteria so AI can support or automate them.
2. **Exceptions:** Define routings and thresholds for human review. AI without an exception path either blocks flow or creates silent errors.
3. **Feedback loop:** Capture outcomes (accepted/rejected predictions, corrected classifications) to improve models and prompts over time.

Where processes vary by product, region or customer tier, either **standardise** or treat each variant as a distinct pathway. Half-standardised processes are a common failure mode.

Sequencing: a practical blueprint

The temptation is to start with the technology.

The discipline is to start with the outcome.

- 1. Pick one important area of the business:** Choose a process where improvement would really matter to the P&L, such as reducing billing errors, speeding up claims handling or improving demand forecasting.
- 2. Define what success looks like:** Be explicit about the measure; fewer errors, faster turnaround, better accuracy. Not just "using AI".
- 3. Stabilise the way the work is done today:** If people are doing the same task five different ways, AI will only amplify the chaos. Standardise the process so it's predictable.
- 4. Get the data in order:** For that process, agree definitions, clean the records and make sure the numbers line up. Good data doesn't mean perfect data everywhere; just good enough for this use case.
- 5. Introduce AI where it can move the dial:** Now layer AI onto the process, classifying, predicting or automating where it makes sense while keeping human oversight where the stakes are high.
- 6. Measure, learn and expand:** Track the business results, not just the model accuracy. Use the lessons from one process to speed up the next.

The goal is to deliver early wins, build confidence and establish a repeatable pattern.

This way, AI adoption becomes cumulative, each success laying the foundation for the next.



Why foundations needn't be enterprise-wide

Enterprise-wide programmes stall under their own weight. Domain-scoped foundations deliver faster impact and build organisational confidence.

Think in **capabilities** at the process level: people, skills, experience, systems, data, and IP applied to a specific outcome. Build that capability deeply in one domain; prove value; codify the patterns (schemas, pipelines, playbooks); and reuse them in adjacent domains.

Over time you will have an enterprise foundation, assembled from proven parts, not designed in a vacuum.

Risk, control and credibility

AI introduces new operational risks. The answer is not to pause, it is to govern intelligently:

- + **Evaluation as a product:** Treat test sets and acceptance thresholds as first-class artefacts.
- + **Security by construction:** Build in access controls and privacy rules from the start.
- + **Explainability appropriate to risk:** Where outcomes affect customers, compliance or cash, prioritise approaches that can be audited.
- + **Human-in-the-loop where it counts:** Route high-risk, low-confidence cases to people; log and learn from corrections.
- + **Vendor pragmatism:** Avoid lock-in by isolating providers behind clear data definitions and contracts.

Credibility with investors and buyers comes from repeatable results, not glossy demos. Foundations create repeatability.

The mid-cap reality: legacy vs digitally native

A digitally native business often has modern systems, coherent data and fewer exceptions. AI can be layered quickly because **capability** in the process sense is already high.

In a 10–20-year-old business, **capability** varies by process: the same workflow may perform differently across teams because the people, skills, systems and data supporting it differ.

Approach these environments with surgical clarity:

- + Pick domains where foundations can be assembled quickly.
- + Retire or bypass legacy integrations where possible.
- + Create one **clean corridor** of process + data + AI that works end-to-end.
- + Expand from that corridor using the same patterns.

This is how you convert a legacy estate from an anchor into an asset, step by step, outcome by outcome.



Conclusion

AI will reshape operational value creation. The strategic question is not whether to adopt it but **how** to do so credibly and profitably within a hold period.

The answer is to **start now, start properly**:

- ✦ **Define the business outcome and bound the scope**
- ✦ **Stabilise the process so work is executed consistently**
- ✦ **Establish domain-level data foundations, clear definitions, ownership accessibility and governance**
- ✦ **Introduce AI to accelerate a process that already works, with oversight and control where the risks matter**
- ✦ **Reuse the patterns to scale across the enterprise**

Foundations first is not delay; it is the shortest path to durable value. AI is different in its potential, but it only pays out when anchored in strong processes and trustworthy data.

Build that capability, one domain at a time, and the technology will do the rest.

Make execution your competitive edge

At **Stratex Plus**, we work with ambitious, PE-backed service businesses to build scalable operating models, embed execution discipline and unlock trapped value.

We don't offer off-the-shelf solutions and will never suggest (someone else's) best practice. We develop and deliver best for you, with focus and precision.

We help you see what's really going on, making the invisible, visible across strategy, operations, systems, and people and design the path to better performance, lower effort and greater investor confidence.

If you're serious about value creation, let's talk.

stratex⁺

Strategy Execution Specialists,
Business Transformation Experts

hello@stratexplus.co.uk
www.stratexplus.co.uk

