



Strategy at the vanguard of execution.

WHITE PAPER • JULY 2026

Evidence before *automation*

Why UK retail and FMCG cannot automate, migrate or invest their way out of a problem they have never measured.

PREPARED FOR

Boards, supply chain, operations and transformation leaders across UK retail and FMCG

PRODUCED BY

STROS Partnerships • STROS Consulting Ltd

Supply Chain • Logistics • Technology • Retail stros.co.uk

Contents

- 1 Executive Summary
 - 2 The Operating Context: UK Retail and FMCG in 2026
 - 3 The Case for Change: Three Converging Deadlines
 - 4 Why Transformation Programmes Underdeliver
 - 5 Process Intelligence: The First Evidence Gap
 - 6 The Second Evidence Gap: Committing Capital Under Uncertainty
 - 7 Where the Value Sits: A Value Map for Retail and FMCG
 - 8 A Practical Approach: From Evidence to Outcome
 - 9 Governance, Risk and the Human Factor
 - 10 The Board-Level Diagnostic
 - 11 Conclusion: The Next Eighteen Months
- About STROS Partnerships

ABOUT THIS PAPER

This white paper is written for leaders who are being asked to deliver measurable operational improvement in conditions that offer very little margin for error. It is deliberately vendor-neutral. It does not advocate a particular platform, and it does not assume that the answer to an operational problem is another system.

Its argument is narrower and, we believe, more useful: that the organisations extracting real value from artificial intelligence are those that first establish an evidence base for how their operations actually work — and that those which skip this step are systematically converting good technology into expensive disappointment.

THE SUMMARY

1 Executive Summary

UK retail and FMCG organisations are operating in the most demanding cost environment in a generation, while simultaneously being asked to modernise their core systems, meet materially expanded regulatory reporting obligations, and demonstrate a return on artificial intelligence investment. These three demands are arriving at the same time, and they are competing for the same scarce internal capacity.

The prevailing response has been to buy capability. Technology budgets across retail and consumer goods are rising, AI is being piloted almost universally, and boards are being presented with roadmaps of considerable ambition. Yet the evidence on outcomes is sobering. Deloitte's 2026 survey of 200 retail and CPG executives found enterprise-wide AI deployment sitting in the single digits in both sectors — organisations are piloting broadly and scaling almost nothing. The constraint is not appetite, and it is rarely the technology itself.

In our experience the constraint is evidential. Most organisations do not have a reliable, shared, current understanding of how their critical operational processes actually run — across brands, sites, markets and business units. Without that foundation, transformation programmes are built on assumption. Automation is applied to workflows nobody has properly examined. ERP migrations carry forward twenty years of accumulated workaround. AI models are trained on process data that describes the exception rather than the rule. The technology performs; the business case does not.

KEY FINDINGS

1. The cost base is now structural, not cyclical. Wage floors, employer National Insurance, business rates and packaging compliance have permanently reset the cost of operating in UK retail and FMCG. Efficiency can no longer be deferred to a better trading year.
2. Three deadlines are converging on 2027. Core platforms are reaching end of life — most visibly SAP ECC, whose mainstream maintenance ends 31 December 2027, though the same process risk attaches to any system migration; packaging EPR fees become recyclability-modulated from the 2026–27 scheme year with nation-of-sale reporting due April 2027; and boards are demanding proof of AI return within the current financial year.
3. Pilots are outpacing production by an order of magnitude. Broad experimentation with negligible enterprise-wide deployment is the defining pattern of AI adoption in both sectors.
4. The binding constraint is evidence, not technology. Organisations cannot automate, migrate or optimise what they cannot accurately describe — nor commit capital confidently to a change they have never tested.
5. There are two evidence gaps, not one. The first concerns how the operation actually works today. The second concerns what a proposed change would actually do to it. Closing one without the other leaves the risk in place.
6. Vendor AI agents stop at the system border. Almost every platform provider now ships agents confined to its own software, while real processes cross four or five systems. Until agents can travel between estates, only the organisation can own the end-to-end view.
7. Evidence is the cheapest insurance available. Establishing a verified baseline and testing options before commitment typically costs a fraction of one per cent of a major ERP or automation programme, and materially de-risks the whole.

This paper sets out the operating context for UK retail and FMCG in 2026, examines why well-funded transformation programmes underdeliver, defines the two evidence gaps that cause it — an unverified

current state and an untested future state — maps where value realistically sits across both sectors, and offers a practical five-stage approach and a board-level diagnostic. It closes with a view on what the next eighteen months require.

THE CONTEXT

2 The Operating Context: UK Retail and FMCG in 2026

Any credible case for operational change must begin with an honest reading of the conditions in which it will be delivered. Those conditions have hardened considerably.

2.1 Retail: demand is fragile and the cost base has been permanently reset

The headline inflation picture has improved — shop price inflation ran at modest levels through the first half of 2026, with food inflation easing to its lowest point in a year — but slower price growth has not translated into lower costs for retailers. The British Retail Consortium has consistently warned that retailers are absorbing cost increases rather than passing them through, and that this absorption has limits.

Three pressures deserve particular attention:

- **Labour.** The National Living Wage rose 4.1% to £12.71 per hour, with sharper increases for younger cohorts, layered on top of earlier employer National Insurance changes. For a labour-intensive store and distribution estate this is a permanent step-change in the cost of every hour worked.
- **Property and rates.** Business rates remain the pressure retailers themselves rank highest, reflecting the fixed-cost burden of large physical estates and continued uncertainty over revaluation.
- **Demand.** Footfall fell 10.7% year on year in April 2026 across all major retail formats according to BRC and Sensormatic data, and consumer behaviour remains value-focused and promotion-dependent. Growth is available, but it is being won rather than received.

The consequences are visible in the corporate record. FTSE retailers issued eight profit warnings in the first half of 2026 — two more than the same period a year earlier — and 2026 became only the third year since 2007 in which the sector recorded more warnings in the second quarter than the first, according to EY-Parthenon. A run of high-profile failures, CVAs and restructurings through late 2025 and early 2026 underlines that the distribution of outcomes has widened. Execution, not strategy, is increasingly what separates outperformers from casualties.

2.2 FMCG: margin compression from both directions

Fast-moving consumer goods businesses face a structurally similar squeeze arriving through different channels. Input, energy and freight costs remain elevated and volatile; tariff exposure, sanctions and foreign exchange movements continue to disrupt planning assumptions; and the retailer relationship has become less forgiving as grocers defend their own margins against discounter competition. Trade terms, service-level penalties and OTIF chargebacks transmit retail pressure directly onto supplier P&Ls.

Layered on top is a regulatory obligation of genuine operational weight. The UK's packaging Extended Producer Responsibility regime shifts the full cost of collecting, sorting and recycling household packaging onto producers — an estimated £1.38 billion annually across the scheme. From the 2026–27 scheme year, fees are no longer flat: they are modulated according to the recyclability of each format, with the first modulated invoices landing in the second half of 2026 and further fee increases signalled for subsequent years. Producers must additionally report nation-of-sale data from 1 April 2027, and prepare for mandatory recyclability labelling thereafter.

WHY EPR IS AN OPERATIONS PROBLEM, NOT A COMPLIANCE PROBLEM

EPR fees are calculated from packaging volume, material and weight at SKU level. Producers whose production and supply records cannot be resolved to that granularity default to site-level reporting and blended maximum fees — paying more than they owe because they cannot evidence what they actually placed on the market.

The British Retail Consortium has reported that the large majority of retailers now face significant administrative burden under the scheme, and that most of the cost is expected to reach consumers. Data quality has moved from a back-office concern to a direct determinant of cost.

2.3 What both sectors now share

Retail and FMCG have converged on a common operating profile. Both are working with permanently higher fixed costs and limited pricing headroom. Both are managing greater process complexity — more channels, more markets, more SKUs, more compliance obligations — with flat or reduced headcount. Both depend on data that must now be accurate at a granularity their systems were never designed to deliver. And both are being asked by their boards to prove, within the current financial year, that technology investment is producing measurable results.

These are the conditions in which change must be delivered. They reward precision and punish assumption.

THE FORCING FUNCTION

3 The Case for Change: Three Converging Deadlines

Cost pressure alone rarely forces structural change; organisations become adept at absorbing it. What makes the present moment different is that three hard deadlines are converging, and each one independently requires an accurate understanding of current-state operations.

Deadline one: the systems clock

Almost every organisation in both sectors is carrying at least one core platform approaching the end of its useful life — an ERP, a warehouse management system, a planning suite, an order management platform, a point-of-sale estate. Replacing any of them is fundamentally the same exercise: lifting a body of processes that nobody has fully documented off one system and re-establishing them on another, while continuing to trade. The scale varies enormously. The failure mode does not.

SAP offers the clearest worked example, because its deadline is external, published and unusually large. Mainstream maintenance for SAP ECC ends on 31 December 2027, with extended maintenance available to 2030 at a significant premium. Full ECC-to-S/4HANA migrations typically run 18 to 36 months. Gartner data indicated that only around 39% of ECC customers had migrated by the end of 2024, and industry commentary through 2026 suggests a substantial proportion of the remainder have not started. The arithmetic is unforgiving: organisations beginning after mid-2026 are compressing a multi-year programme into a window that will not accommodate it, while competing for specialist talent that is already scarce and becoming more expensive.

For retail and FMCG, a change of this kind is not an IT milestone. These systems are where demand planning, replenishment, trade promotion, order-to-cash and financial control actually live. A migration executed without a verified picture of current-state process reproduces existing inefficiency on a more expensive platform — and does so under time pressure.

Organisations not running SAP should not read the above as exemption. The relevant point is the opposite. They face the same process exposure without the discipline that an external deadline imposes. A warehouse management replacement, a planning system upgrade, a new order management layer or a first significant automation deployment carries identical risk, but is timed by internal appetite and vendor pressure rather than a published end-of-support date — which in practice means it is scoped faster, scrutinised less, and justified on the same unverified assumptions. A hard deadline is the one advantage SAP customers have that everyone else does not.

Deadline two: the regulatory data clock

Modulated EPR fees are already in effect for the current scheme year; nation-of-sale reporting follows in April 2027. Both require SKU-level packaging and distribution data that most organisations cannot currently produce without significant manual reconciliation. The organisations that will pay the least are those that can evidence the most — and evidencing requires knowing precisely how packaging specification, procurement, production and distribution data flows through the business today.

Deadline three: the board's patience

The AI conversation has moved decisively from experimentation to accountability. Industry surveys through 2026 show near-universal engagement — the overwhelming majority of retail and CPG organisations are actively using or assessing AI, and around nine in ten intend to increase AI budgets — alongside rapidly growing interest in agentic systems that execute rather than merely recommend. But the same research shows enterprise-wide deployment stuck in single digits, and McKinsey has put the proportion of organisations genuinely scaling agentic AI at under a quarter.

The gap between those figures is where credibility is lost. A third consecutive year of pilots without production deployment will not survive contact with a finance director defending a compressed budget.

THE COMMON DENOMINATOR

A systems migration, an EPR reporting obligation and an AI deployment appear to be three unrelated programmes owned by three different functions. They are not. Each one fails in the same way and for the same reason: an incomplete, contested or out-of-date understanding of how the underlying process actually works.

Establishing that understanding once, properly, serves all three. Treating them as separate initiatives means paying for the same discovery work three times — and getting three inconsistent answers.

THE FAILURE MODE

4 Why Transformation Programmes Underdeliver

Transformation programmes in retail and FMCG rarely fail because the technology does not work. They underdeliver for a small number of recurring, structural reasons.

4.1 Process variance is invisible until it is expensive

In multi-brand, multi-market or acquisitive organisations, the same nominal process — goods receipt, promotional set-up, supplier onboarding, returns — typically exists in several materially different versions. Each was a reasonable local adaptation at the time. Collectively they make enterprise-level design,

automation and reporting extremely difficult, because there is no single process to design for. Variance is usually discovered during build or, worse, during user acceptance testing, at the point when change is most expensive.

4.2 Documentation describes intention, not reality

Most organisations have process documentation. Very little of it describes what people actually do. Standard operating procedures capture the designed path; operational reality is dominated by exceptions, workarounds, informal escalations and the accumulated judgement of experienced staff. When automation is built from documentation, it automates the designed path and breaks on everything else — which is where most of the volume and nearly all of the cost sits.

4.3 Critical knowledge is held by individuals, not systems

A significant proportion of operational capability in both sectors resides in the heads of long-serving people. In a market where labour costs are rising and headcount is under review, this is a live risk. Every unplanned departure removes process knowledge that was never captured, and each one degrades the organisation's ability to change safely.

4.4 Spreadsheets and email are load-bearing

Between the major systems sits a substantial informal layer of spreadsheets, shared mailboxes, messaging threads and manual reconciliation. This layer is where a great deal of real work happens, and it is almost entirely invisible to system-based analysis. Any current-state assessment that draws only on system logs will systematically understate both the effort involved and the risk carried.

4.5 Business cases are built on assumption

Where the current state has not been measured, benefits are estimated from vendor benchmarks and analyst averages rather than from the organisation's own baseline. The programme is then held to a target that was never grounded in its own operations. When benefits fail to appear, the resulting dispute is unresolvable — because nobody can establish what the starting position was.

THE PATTERN IN ONE SENTENCE

Organisations invest heavily in changing operations they have never accurately measured, and are then surprised when the change delivers less than forecast.

THE FIRST GAP

5 Process Intelligence: The First Evidence Gap

Process intelligence is the discipline of establishing an evidence-based, current-state understanding of how work actually happens across people, process and systems — and maintaining it as a living asset rather than a one-off study.

It is worth being precise about what this does and does not mean, because the term is used loosely.

WHAT PROCESS INTELLIGENCE IS	WHAT IT IS NOT
A verified account of how work is actually performed, including exceptions, hand-offs and decision points	A refresh of existing process documentation
A reconciliation of process variants across brands, sites and markets into a single governed view	A mandate to standardise everything regardless of local commercial logic
A quantified baseline against which benefit claims can be tested	A benchmark comparison against other organisations
Coverage of the informal layer — spreadsheets, email, manual intervention — alongside system activity	System log analysis alone
A prioritised, costed view of where automation and AI will actually pay	A recommendation to replace core systems
A durable asset that supports ERP migration, compliance reporting and AI deployment simultaneously	A deliverable that expires when the consultants leave

5.1 Why it matters more now than it did three years ago

Two things have changed. First, the tolerance for unproven investment has collapsed; every significant programme now competes for scarce capital against a demanding hurdle rate. Second, the technology itself has become dependent on process quality in a way that earlier automation was not.

Rule-based automation failed loudly and locally: a robot broke, someone fixed it. Agentic AI — systems that reason across sources and execute rather than recommend — fails quietly and systemically. An agent operating on an inaccurate model of a process will act confidently on the wrong basis, at scale, and its errors will be difficult to detect until they surface in inventory, service levels or the ledger. As industry commentary has increasingly observed, most cancelled agentic pilots do not fail on model quality. They fail on the estate that was never properly connected, on terms like “lead time” that mean four different things in four different systems, and on constraints the planners knew about but nobody wrote down.

Process intelligence is what closes that gap. It is also, in cost terms, the cheapest component of any major programme — typically a small fraction of one per cent of an ERP migration budget — while materially reducing the risk carried by the remainder.

5.2 Passport control: the agent that cannot cross the border

There is a second reason the evidence has to be owned at enterprise level, and it is becoming more acute with every product release. Almost every major platform provider now ships agentic capability inside its own software. The ERP vendor has agents. So does the warehouse management system, the planning suite, the transport platform, the finance system, the e-commerce stack and the point-of-sale estate. Each is genuinely capable, and each is confined to the territory its vendor controls.

Very few retail or FMCG businesses run a single vendor’s suite end to end. The typical estate is assembled rather than designed: an ERP from one provider, warehouse management from another, planning from a third, transport and returns from others again, joined by integration middleware and a substantial layer of manual reconciliation. Each vendor’s agent holds citizenship of its own system. Inside those borders it moves freely. At the edge it reaches passport control — and it has no papers.

The difficulty is that business processes are not citizens of any system. A single order-to-cash cycle, promotional launch or supplier claim crosses four or five platforms and several human hand-offs before it completes. The process is the whole journey. The agent is confined to one leg of it, optimising against the

only metrics it can see — which are its own. An agent tuned to despatch efficiency in the warehouse cannot see the delivery promise made at the point of sale, and will happily improve its own number while degrading the outcome that actually matters.

P A S S P O R T C O N T R O L

Vendor agents are issued citizenship of a single system. They can act freely within it and not at all beyond it.

Business processes hold no citizenship anywhere. They cross every border in the estate, and it is at those borders — the hand-offs, the reconciliations, the queues between systems — that most of the cost and nearly all of the delay accumulate.

Until agents can be issued passports and a right of travel, the organisation itself has to act as the customs union: deciding what crosses, in what form, carrying what context, and under whose authority. That is an enterprise responsibility, and it cannot be delegated to a software provider.

Three practical consequences follow.

- **No vendor can supply the end-to-end map.** Each provider can describe its own territory accurately and none can describe the route. Only the organisation is positioned to hold the whole journey — which is precisely what an enterprise-owned process evidence base is for. Where that map does not exist, the estate is effectively being designed by whichever vendor is most persuasive in the room.
- **Interoperability belongs in procurement, not in year three.** The questions to ask before signing are specific: what can this agent see, what can it act on, where does its authority stop, what context can it hand to an agent outside this estate, and which interoperability standards does it support today rather than on a roadmap. Emerging protocols for passing context between systems are promising, but they do not resolve a twenty-year-old integration layer on their own.
- **The border is where the value is.** Because every vendor optimises inside its own boundary, the space between systems is the least automated and most manually reconciled part of the chain — and no single provider is commercially incentivised to fix it. That makes cross-system hand-offs both the largest remaining pool of operational value and the one that only the organisation itself can go after.

This is not an argument against vendor agents, which will often be the fastest route to value inside a given domain. It is an argument for knowing where the borders fall before deploying them — and for holding an end-to-end process view that no supplier is able to give you.

It addresses, however, only half of the evidence problem. Knowing precisely how an operation works today does not, by itself, tell you whether a proposed change to it will work. That is the second gap, and it is where the largest sums of money are at stake.

T H E S E C O N D G A P

6 The Second Evidence Gap: Committing Capital Under Uncertainty

The decisions that most affect operational performance in retail and FMCG are not incremental. They are step-change commitments: consolidating a distribution network, automating a site, replacing materials handling equipment, redesigning a labour model, adding a channel or format, or standing up capacity for a peak that has not happened yet. These are the largest discretionary sums an operations board will approve, they are difficult to reverse, and they carry consequences for years.

They are also, in many organisations, the decisions supported by the weakest evidence.

6.1 Step-change decisions are being made on linear assumptions

The conventional basis for a major operational investment is a static model: a spreadsheet built on average volumes, vendor-supplied throughput rates, assumed labour productivity and a straight-line ramp to steady state. It is a linear representation of an environment that is not linear. Real operations are governed by interaction effects — congestion at a pinch point, queueing behind an exception, labour availability that varies by shift and by season, the way a single delayed inbound cascades across a pick face for the rest of the day. None of this appears in a spreadsheet, and all of it determines whether the projected benefit materialises.

The result is a recognisable pattern. A business case is approved on a modelled throughput figure. The installed solution performs to specification in isolation, but the surrounding operation does not behave as the model assumed. Benefits arrive late, partially, or not at all — and because no one modelled the alternatives, the organisation cannot establish whether a different configuration would have performed better.

6.2 The cost of automation false starts

A decade of partially realised automation value has left many boards understandably cautious about capital-intensive operational change. That caution is rational at the level of the individual decision and damaging in aggregate: productivity stagnates while cost pressure continues to build. Faced with a choice between defending the balance sheet and pursuing productivity, organisations have tended to defend — and in the current cost environment, standing still is no longer a neutral position.

The way out of that impasse is not greater boldness. It is better evidence about consequences before capital is committed.

6.3 Testing the future state before buying it

Simulation and digital twin capability has matured to the point where it is a practical decision-support tool rather than a specialist engineering exercise. A validated model of an operation — built on the process evidence described in Section 5 — allows options to be run against realistic conditions before any commitment is made: peak as well as average, disruption as well as steady state, several candidate designs side by side rather than one preferred option defended after the fact.

The value is not the model. It is the ability to make trade-offs explicit and quantified. Cost-to-serve against service level. Throughput against capital outlay. Resilience against labour exposure. Presented in that form, a capital paper stops being an argument about whose assumptions are more credible and becomes a comparison of tested alternatives — which is a materially different conversation to have with a board or an investment committee.

THE TWO EVIDENCE GAPS

Gap one — current state. We cannot accurately describe how our operation works today. Consequence: we automate, migrate and report on a process we have not verified.

Gap two — future state. We cannot test what a proposed change would actually do. Consequence: we commit capital on linear assumptions and discover the interaction effects after installation.

The two are connected. A simulation is only as reliable as the process evidence beneath it — which is why current-state discovery is the prerequisite for future-state testing, not an alternative to it.

6.4 Decisions worth testing before committing

DECISION	THE QUESTION A TESTED MODEL ANSWERS
Network footprint	What does consolidating, relocating or adding a site do to cost-to-serve, service level and resilience under real demand patterns?
Automation and robotics	At what volume does the investment actually pay, how does it behave at peak, and what happens to the operation when it is unavailable?
Materials handling and storage	Does the proposed configuration hold up against the SKU profile and order mix we will have in three years, not the one we have now?
Labour model and shift design	Where is the operation constrained by people rather than equipment, and which shift patterns absorb volatility at least cost?
Peak readiness	What breaks first as volume rises, and what is the cheapest intervention that moves the breaking point?
New channel or format	What does adding this flow do to the existing operation — and can the current estate absorb it without degrading service?

Not every decision warrants this treatment. The test is straightforward: where the commitment is large, difficult to reverse, and dependent on how several variables interact, the cost of modelling it properly is trivial against the cost of getting it wrong. Where it is small, reversible and well understood, it is not.

WHERE VALUE SITS

7 Where the Value Sits: A Value Map for Retail and FMCG

The following map identifies the operational domains in which evidence-led process improvement most reliably produces measurable results in each sector. It is intended as a prioritisation aid rather than an exhaustive inventory.

DOMAIN	RETAIL APPLICATION	FMCG APPLICATION	INDICATORS TO BASELINE
Demand & forecasting	Store and channel-level forecast accuracy; promotional uplift modelling	Customer-collaborative forecasting; POS and syndicated data reconciliation	Forecast error, bias, planner touch rate
Inventory & availability	Allocation and replenishment across stores and DCs; markdown exposure	Finished goods and raw material cover; obsolescence and write-off	On-shelf availability, stock cover, write-off value
Promotions & trade	Promotion set-up, execution and post-event evaluation	Trade promotion management, deductions and claims settlement	Set-up lead time, claim cycle time, deduction leakage
Supplier & customer	Supplier onboarding, compliance and performance management	Service level performance, OTIF and chargeback exposure	OTIF, chargeback value, dispute resolution time
Order-to-cash	Multi-channel order orchestration and returns settlement	Invoice accuracy, dispute and credit note volume	Days sales outstanding, credit note ratio
Warehouse & logistics	Goods-in, pick accuracy, store delivery scheduling	Production-to-despatch flow, co-packing and third-party sites	Cost per case, dock-to-stock, labour hours per unit
Returns & reverse flow	Online returns processing, refurbishment and disposition	Recall readiness, traceability and reverse logistics	Return cycle time, recovery rate
Compliance & reporting	EPR and packaging data across own-brand ranges	SKU-level packaging data, nation-of-sale reporting, PPT exposure	Data completeness, fee band accuracy, manual effort
Systems change	Current-state evidence for migration scope and design — ERP, WMS, OMS or POS	Consolidation of process variants across sites and markets	Variant count, custom object usage, test coverage
Network & capital design	Store fulfilment models, DC consolidation, automation business cases	Site footprint, co-packing capacity, automation and MHE investment	Cost-to-serve, throughput at peak, payback confidence

7.1 A note on benefit ranges

Analyst estimates for AI-enabled supply chain improvement commonly cite reductions in inventory carrying cost in the region of 10–30%, and improvements in planning productivity of 20–50%. Vendor-commissioned research has attributed multi-million-pound three-year benefits to forecasting and inventory optimisation deployments in retail and consumer goods.

These figures are useful for framing ambition. They are not a business case. A defensible business case requires the organisation's own baseline, measured before the intervention, in the specific domains where it intends to act. Establishing that baseline is the first practical step, and it is one of the few steps that pays for itself regardless of what is decided afterwards.

HOW WE WORK

8 A Practical Approach: From Evidence to Outcome

The approach below is deliberately sequential. Each stage produces a decision-grade output, and each is designed so that an organisation can stop, act on what it has learned, and re-enter later without wasted effort.

STAGE	OBJECTIVE	METHOD	OUTPUT
One	Discovery	Structured interviews, operational walkthroughs, documentation review and system evidence across the selected domains	Verified current-state process map, including exceptions and the informal layer
Two	Insight	Quantification of variance, value leakage, bottlenecks and manual effort against measured volumes	Prioritised opportunity register with sized impact and confidence rating
Three	Transformation planning	Sequencing by value, feasibility and dependency; alignment to ERP and compliance timelines	Costed roadmap with named owners and measurable outcome targets
Four	Automation & governance	Deployment of agreed improvements with controls, audit trail and human-in-the-loop escalation defined in advance	Live improvements operating under explicit governance
Five	Continuous improvement	Measurement against the original baseline; refresh of the process evidence base as operations change	Sustained performance and a maintained process asset

9.1 Sequencing in a compressed window

Where a platform migration is already in view, the sequence does not change but the scope discipline tightens. Discovery should be confined initially to the processes that will determine migration scope and design — typically order-to-cash, procure-to-pay, plan-to-produce and the reporting obligations attached to them. Broader coverage can follow once the migration path is set. The objective at this stage is not completeness; it is ensuring that irreversible design decisions are made on evidence rather than on the recollection of whoever was in the room.

9.2 The ninety-day question

Boards reasonably ask what can be demonstrated within a quarter. In our experience a well-scoped discovery and insight phase across two or three high-value domains can deliver, within ninety days, a verified current-state baseline, a quantified opportunity register, and at least one improvement in live operation. That is a realistic expectation. A promise of enterprise-wide transformation in the same window is not, and boards are increasingly able to tell the difference.

CONTROLS

9 Governance, Risk and the Human Factor

9.1 Governance must precede autonomy

As systems move from recommending to executing, governance ceases to be a documentation exercise. Decision rights, escalation thresholds, auditability and the boundary of autonomous action need to be defined before deployment, not retrofitted after an incident. Where several vendor agents operate across a single process, that definition has to be made once at enterprise level rather than separately inside each supplier's console, or the organisation ends up with several agents acting confidently under different assumptions about the same order. Organisations seeing the strongest returns are consistently those that have embedded human-in-the-loop controls and explainability from the outset — not because it is prudent in the abstract, but because it accelerates adoption by the people whose judgement the system depends on.

9.2 Name an owner

Programmes of this kind require a single accountable owner with authority over both the process and the technology decision. Where ownership is split between a transformation function that owns the tooling and an operational function that owns the outcome, neither can be held to the result. The most effective structure places accountability with the P&L owner who will carry the benefit.

9.3 The workforce dimension

Process discovery is experienced by operational teams as scrutiny unless it is framed and conducted otherwise. The people who know how the process really works are also the people with most to lose from an inaccurate account of it, and they will calibrate their candour accordingly. Programmes that treat frontline and planning staff as the primary source of evidence — and that visibly act on what those teams report — obtain materially better information than those that treat them as subjects of the exercise. This is not a soft consideration; it is the principal determinant of data quality in stage one.

9.4 Risks worth naming explicitly

- **Scope drift.** Discovery expands to cover processes with no near-term decision attached to them, consuming budget without changing any outcome.
- **Analysis without action.** A high-quality baseline is produced and then filed. The evidence base has no value independent of the decisions it informs.
- **Standardisation for its own sake.** Process variance is sometimes commercially justified. The objective is governed variation, not uniformity.
- **Automating a bad process faster.** The most common and most expensive failure mode, and the one this approach exists to prevent.

THE QUESTIONS

10 The Board-Level Diagnostic

The following questions are intended for use in a board or executive committee setting. They are diagnostic rather than rhetorical: an organisation that can answer every one of them with evidence is well positioned, and one that cannot has identified its own starting point.

1. Can we produce, today, an accurate description of how our five most critical operational processes actually run — and would our frontline teams recognise it?
2. How many versions of each of those processes exist across our brands, sites and markets, and which of those variants are commercially justified?
3. What proportion of the work in those processes is completed outside our core systems, in spreadsheets, email or manual reconciliation?
4. Is the scope of our next major system change — ERP, warehouse management, planning or otherwise — built on measured current-state evidence, or on documentation and recollection?
5. Can we evidence our packaging and product data to the granularity that modulated EPR fees and nation-of-sale reporting require — or are we paying blended rates because we cannot?
6. When we last approved a major operational capital investment — an automation, a site consolidation, a network change — which alternatives were modelled, and against what conditions?
7. Do we know what breaks first in our operation as volume rises, and what the least expensive intervention would be that moves that point?
8. For each AI initiative under way, what baseline was measured before it started, and who is accountable for the delta?
9. How many of our AI use cases have reached enterprise-wide production, as distinct from pilot or proof of concept?
10. For each vendor-supplied AI agent we have deployed or are considering, where exactly does its authority stop — and who owns the process on either side of that boundary?
11. Where an AI or automated system takes an operational decision, who has authority to override it, and is that decision auditable?
12. What operational knowledge would leave the business with our ten longest-serving operational staff, and where is it captured?
13. If we were asked to demonstrate a measurable operational outcome within ninety days, what would we choose — and could we prove the improvement?

THE THESIS

11 Conclusion: The Next Eighteen Months

The next eighteen months will be unusually consequential for UK retail and FMCG operations. For organisations running SAP the migration window closes at the end of 2027, and effectively closed for late starters in the middle of this year; for everyone else the same process risk sits inside whatever platform change is already on the roadmap, without a published date to enforce discipline. Modulated packaging fees are already being invoiced, with further increases signalled. Consumer demand is fragile and the cost base will not revert. And the period of tolerance for AI investment without demonstrated return is ending.

At the same time, the capital decisions that will determine operational performance for the rest of the decade — network design, automation, labour models — are being taken now, in conditions where the margin for error is smaller than it has been for years.

None of this argues for more technology. It argues for better evidence — and for the discipline to establish that evidence before, rather than after, committing capital to change. The organisations that will emerge from this period in the strongest position are unlikely to be those that adopted AI earliest. They will be those that understood their own operations most precisely, and could therefore act on them with confidence.

The future of retail and FMCG will not be defined by AI adoption. It will be defined by how effectively organisations use AI to understand, improve and govern the way work actually gets done.

THE SINGLE RECOMMENDATION

Select two or three operational domains where the cost of being wrong is highest — typically those attached to a live systems migration, compliance or availability commitment. Establish a verified, quantified current-state baseline across them within ninety days. Then decide what to automate.

Where a step-change capital commitment is already in view, test the alternatives against realistic conditions before the paper goes to the board, not after the equipment is installed.

The sequence matters more than the speed. Evidence before automation.

BY INTRODUCTION


About STROS Partnerships

STROS is a new kind of consultancy, sitting at the high-stakes intersection of supply chain, technology, logistics and retail — anchoring visionary strategy to elite execution. One firm, two vehicles: STROS Consulting delivers strategy, advisory and C-suite leadership; STROS Partnerships deploys network, people and infrastructure directly into the growth engines of high-potential platforms.

The firm carries the direct experience of its founder, Ross Eggleton FCILT — formerly Executive Group Director of Logistics, Supply Chain & Technology and Group Manufacturing COO at Morrisons, and VP of Business Transformation at GXO. That record of driving operational excellence, cost transformation and durable change inside complex environments is the foundation everything at STROS is built on.

We are sector-first rather than platform-first. Our starting point is the operating problem, not a preferred technology — and we work with a deliberately small circle.

To learn more, visit stros.co.uk



This document is provided for general information. It does not constitute professional advice and should not be relied upon as such. Third-party research and market data referenced in this paper are attributed to their originating organisations and were current at the time of writing. STROS Consulting Ltd accepts no liability for decisions taken on the basis of this document alone.

© 2026 STROS Consulting Ltd, trading as STROS Partnerships. Registered in England and Wales, company number 12214146. All rights reserved.