LOCAL GOVERNMENT AI TRANSFORMATION:

FROM STRATEGY TO VALUE

Connecting the SMART: Maturity Model, SMART: AI TOM, and the SMART: Unified AI Platform to deliver real, measurable, sustainable value from AI

A White Paper by

Martin Neale – CEO, ICS.AI

Updated V2.1 edition - March 2026

ICS.AI

WHITE PAPER OVERVIEW

AI is not simply another software tool. It changes how services are designed, how knowledge is structured, how work is delivered, and how digital systems themselves are produced. For councils and other trusted organisations, the central challenge is therefore not experimentation, but governed scale: how to move from scattered pilots to an AI-Native operating model that can create value safely, consistently, and at portfolio level.

1. EXECUTIVE SUMMARY

Public sector and other trusted or regulated organisations are under sustained pressure to deliver multi-million-pound in-year savings while maintaining (or improving) service quality. The National Audit Office has highlighted that pressures on local government finances and rising demand continue to strain sustainability [1].

The financial opportunity is substantial and now well-evidenced. Conservative modelling based on The Ministry of Housing, Communities & Local Government Revenue Outturn 2024-25 data and independent research indicates that a typical unitary authority can **achieve £8m - £19m per year** in combined income generation and efficiency savings [24]. These figures are based on lower-bound assumptions and are achievable only through disciplined implementation.

Council Type	Income Generation	Efficiency Savings	Combined
District	£0.6m - £1.4m	£0.9m - £1.4m	**£1.5m - £2.8m**
Unitary	£3.5m - £8.0m	£4.4m - £10.9m	**£7.9m - £18.9m**
County	£6.9m - £15.3m	£8.9m - £24.1m	**£15.8m - £39.4m**

The context has now shifted. Agentic AI systems are changing the economics of software creation and operation. Capabilities that once required long project timelines and large specialist teams can increasingly be delivered through AI-assisted and AI-orchestrated production systems. In trusted sectors, however, uncontrolled autonomy is unacceptable. The challenge is not access to AI tools; it is creating the governed infrastructure that allows organisations to use AI safely at scale.

The core question is no longer 'can AI help?' but **'how do we capture this value predictably, at scale, and safely?'**

Adoption is now widespread - 78% of organisations reported using AI in 2024, rising to 88% in McKinsey's 2025 global survey [2][3] – yet enterprise-level value remains elusive. BCG reports that 74% of organisations have yet to demonstrate tangible value from AI [4], and Gartner predicts that 30% of generative AI projects will be abandoned after proof of concept by end of 2025 [5].

The blocker is not technical curiosity or lack of use cases. It is **operating-model failure.** Most organisations get stuck in the AI-Curious stage because they introduce tools without redesigning service delivery, workforce practices, knowledge architecture, governance, and value measurement for an AI-enabled environment.

This paper connects three SMART: elements into a single logic chain - the **SMART: AI Maturity Model**, the **SMART: AI Target Operating Model (AI TOM)**, and the **SMART: Unified AI Platform** - to show how organisations can move from experimentation to measurable, sustainable value.

The central thesis is simple: organisations **cannot achieve maximum or predictable value from AI without a joined-up, organisation-wide, Senior Leadership Team (SLT)-driven approach**. Material outcomes require material change: investment, standardisation, training, governance, and a portfolio of tens to hundreds of use cases that touch every team and (where relevant) every service user.

Start with a senior leadership decision on ambition. In practice there are two targets that matter: AI-Native and AI-Transcendent. AI-Native delivers significant returns within conventional business cycles but still requires an AI-TOM and AI for All. AI-Transcendent adds a cadence shift: faster strategy cycles, more frequent standards updates, stronger optimisation loops, and agentic automation as a core organisational capability.

2. THE STRUCTURAL SHIFT: FROM STATIC SOFTWARE TO GOVERNED AGENTIC SYSTEMS

For more than thirty years enterprise software followed a stable pattern. Applications were designed by human architects, built by development teams, delivered through long project cycles, and then maintained through periodic upgrades. Over time, digital estates accumulated layers of systems that were expensive to build, slow to change, and increasingly difficult to govern as organisational requirements evolved.

That model is now changing. Agentic AI systems – often implemented as coordinated swarms of specialised agents – can analyse existing estates, design solution architectures, generate production code, structure organisational knowledge, test systems, and deploy improvements. What previously took months can increasingly be executed in days or weeks.

More importantly, software no longer needs to be treated as a static asset. It can be continuously created, improved, and reconfigured by AI-enabled production infrastructure. This creates enormous strategic upside, but it also introduces new risks: incorrect policy interpretation, inaccessible or non-compliant services, inconsistent outputs, and behaviour that becomes difficult to explain or control.

For regulated environments such as local government, healthcare, and education, those risks are unacceptable. The issue is not whether the capability exists; it clearly does. The issue is whether the governance infrastructure exists to apply that capability safely. The answer is the combination of a structured transformation programme and a platform architecture that embeds control directly into the way AI is deployed.

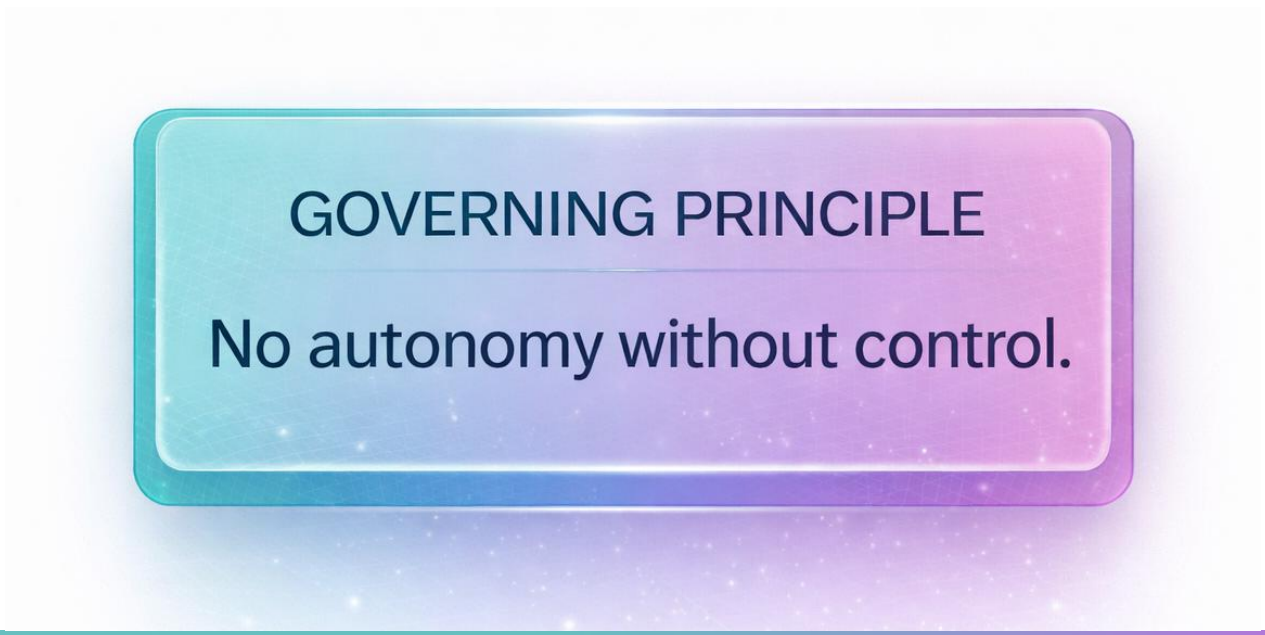

3. THE VALUE IMPERATIVE: WHY INCREMENTAL APPROACHES FAIL

Most organisations start with pilots: a few enthusiastic teams, a small set of use cases, and one or two tools. Pilots are useful for learning, but they systematically fail to deliver material, predictable outcomes because they do not address the scale, integration, and governance requirements of real-world delivery.

3.1 The Evidence: Adoption Is High, but Value Realisation Lags

Multiple independent datasets converge on the same diagnosis: AI activity is high; scalable material value is the exception; pilots often fail because operating foundations aren't in place.

Global Adoption (High and Rising)

Finding	Source
78% of organisations using AI in 2024 (up from 55% in 2023)	Stanford AI Index 2025 [2]
88% regularly using AI in at least one business function	McKinsey Global Survey 2025 [3]
71% regularly using GenAI specifically	McKinsey Global Survey 2025 [3]
42% of enterprises (1,000+ employees) have actively deployed AI; another 40% exploring	IBM Global AI Adoption Index [7]

The Value Gap (Material Impact Remains Elusive)

Finding	Source
Only 31% of organisations have begun to scale AI enterprise-wide	McKinsey Global Survey 2025 [3]
74% of companies have yet to show tangible value from AI	BCG, AI Adoption in 2024 [4]
Only 26% have capabilities to move beyond Proof of Concepts (PoCs) and generate tangible value	BCG, AI Adoption in 2024 [4]
Only 23% can tie GenAI initiatives to new revenue or lower costs	Bain, GenAI Value Realisation Survey [6]
≥30% of GenAI projects will be abandoned after PoCs by end 2025	Gartner, GenAI Project Abandonment Forecast [5]

Key takeaway: Adoption is no longer the core problem. The problem is converting adoption into material value through repeatable operating mechanisms.

3.2 UK Public Sector: The Same Pattern, with Extra Constraints

Central Government (UK National Audit Office (NAO))

The UK National Audit Office provides concrete 'state of play' numbers [8]:

- 37% of government bodies had deployed AI; another 37% were actively piloting or planning
- **Only 21%** reported having an AI strategy (while 61% had plans to develop one)
- **70%** said skills were a barrier to AI adoption

Local Government (LGA State of the Sector)

The LGA's 2024 baseline survey (74 responding councils) and 2025 update (104 councils) shows adoption is rising fast, but maturity and benefits lag [9]. By 2025, 95% of councils are using or exploring AI, yet only 22% report realised cost savings - and 56% don't know whether they've achieved any [9]. This is not a technology problem; it's a measurement and value-realisation problem.

Self-Assessed Readiness Across Core Dimensions:

Readiness Dimension	'Very/Fairly Ready'
Technology	53%
Data	28%
Workforce (skills/knowledge/expertise)	20%

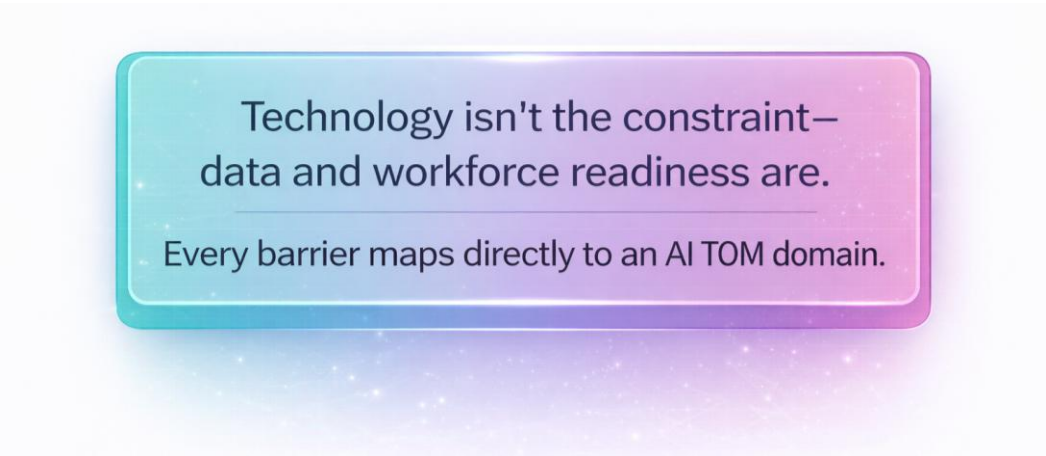

3.3 Why Bottom-Up Adoption Fails

Fragmented, bottom-up adoption tends to fail for five recurring reasons:

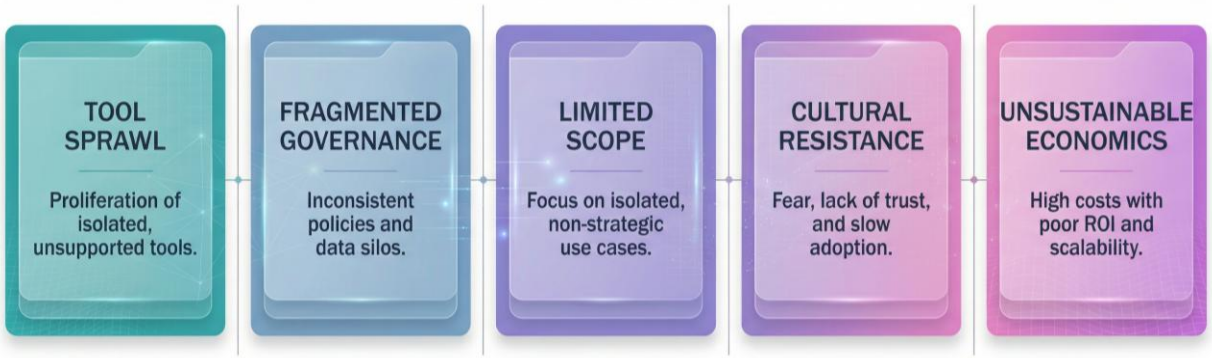

- **Fragmented investment:** multiple teams procure different tools, duplicating costs, fragmenting governance and creating inconsistent user experience.
- **Limited scope:** department-level use cases cannot fix end-to-end processes or cross-cutting demand.
- **Governance complexity:** each tool introduces new risk, assurance, and compliance work; the overhead compounds as you scale.
- **Cultural resistance:** without visible SLT mandate and modelling, adoption remains optional and uneven.
- **Unsustainable economics:** premium per-seat subscriptions often make true 'AI for All' financially unrealistic.

3.4 A Simple Value Equation

$$\text{Annual AI Value} \approx (\text{Coverage of work}) \times (\text{Adoption}) \times (\text{Productivity lift per interaction}) \times (\text{Quality and safety}) - (\text{Total cost to run the capability})$$

If coverage is limited to a small subset of roles or a single channel, the value ceiling remains low regardless of how impressive the pilot appears. The only durable way to raise the ceiling is an organisation-wide approach: common platform, common patterns, common governance, and universal access.

4. THE SCALE REQUIREMENT: WHY 'AI FOR ALL' IS NOT OPTIONAL

Organisations seeking multi-million-pound in-year savings rarely achieve them through a handful of isolated use cases. In practice, material outcomes usually require tens to hundreds of AI-enabled workflows spanning service users, back-office processes, and everyday staff work. AI must therefore become infrastructure, not an innovation side project.

4.1 The Productivity Potential Is Real

Evidence from UK government pilots shows significant productivity potential:

- **Cross-government GenAI trial:** Using GenAI tools saved users 26 minutes per day on average, amounting to nearly two weeks per year per person [10].
- **NHS Microsoft 365 Copilot trial:** Reported potential time savings of 400,000 hours per month for NHS staff [11].

These are strong proofs of *productivity potential* - but the TOM question is: **how do we convert minutes saved into service performance, budget impact, or capacity release?** Without an operating model that connects individual productivity to organisational outcomes, the value remains theoretical.

Without a sanctioned, organisation-wide platform, these individual productivity gains are captured unevenly, governed inconsistently, and often realised through tools the organisation cannot see or control.

4.2 Three Implications for Scale

Scaling AI is not simply a technology change. It is a behavioural and operating-model challenge. In practice, three implications to scale follow.

1. **Every staff member** needs access to approved AI tools for their role, backed by training and safe ways of working – or they will source their own.
2. **Service users** (citizens, patients, students, customers) need consistent, high-quality AI-enabled access through the channels they already use.
3. **AI agents and automation** must be deployable at scale without bespoke integration and governance every time.

In practice, a level of AI adoption already exists on a continuum across the workforce. Some staff are cautious and compliant, others are actively experimenting, and a growing minority are already using public or unsanctioned tools to improve productivity - this behaviour is typically rational rather than malicious.

If organisations do not provide a sanctioned, usable alternative (i.e. an organisation-approved AI platform with defined governance and controls), people will route around constraints. At scale, this creates unmanaged risk, fragmented cost, inconsistent outcomes, and governance blind spots. Providing an organisation-wide, approved AI platform therefore becomes a prerequisite for safe, affordable scale, not an optional control measure.

This is why a single, top-down, organisation-wide programme is typically the only viable delivery pattern.

4.3 The three value domains that must scale together

Domain	What It Is	Why It Matters for Value
1. Front Door	AI that handles enquiries across channels - phone, webchat, email, SMS, social, and assisted digital - with safe escalation.	Reduces avoidable demand, improves access, and shifts human capacity to complex cases.
2. Workforce AI	AI assistance embedded in everyday staff work: drafting, analysis, policy guidance, case summarisation, and knowledge support.	Compounds productivity across every team, improves consistency, and displaces unsanctioned shadow AI.
3. Agentic Back Office	AI agents that execute multi-step workflows across systems, with rules, logging, and oversight.	Delivers end-to-end cycle-time reduction and cashable efficiency.

5. THE SMART: AI TRANSFORMATION PROGRAMME - FROM AI-CURIOUS TO AI-NATIVE

Most organisations today still operate using a model designed for the pre-AI era. Services are accessed through websites and forms; staff navigate complex internal systems; knowledge sits in documents and web pages; and governance lives in policy documents rather than in runtime controls. AI-Native organisations operate differently.

- Citizens interact through conversational service journeys rather than fragmented channel hand-offs.
- Staff work alongside AI in daily workflows rather than treating AI as an occasional add-on.
- Operational processes are orchestrated by AI agents where it is safe and valuable to do so.
- Organisational knowledge is structured for machine use, not only for human reading.
- Governance is embedded directly into digital infrastructure.

The SMART: AI Transformation Programme is the structured methodology for moving organisations towards this model. It recognises that transformation cannot occur simply by deploying AI tools. Organisations must redesign their operating model, governance model, and knowledge architecture for an AI-enabled environment.

5.1 Four Individual Maturity Stages (How People Work)

Level	Characteristics	Typical Behaviours
1. AI-Agnostic	AI is discussed, but there is little formal deployment.	No standard tools, reactive governance, scattered personal usage.
2. AI-Curious	Pilots and pockets of deployment emerge.	Tool sprawl, inconsistent governance, local value rather than enterprise value.
3. AI-Native	AI is embedded into roles and services. AI-TOM in place; AI for All; governed delivery.	Unified platform patterns, measurable outcomes within conventional business cycles.
4. AI-Transcendent	AI-first operating cadence and continuous optimisation become normal.	Monthly strategy cycles, daily monitoring, weekly changes, agents as a core capability.

5.2 Organisational Maturity (How the Organisation Operates)

While individual capability determines how people work, organisational maturity determines whether AI can be scaled, governed, and translated into sustained value. Most organisations experimenting with AI do not fail because the tools are incapable. They fail because the organisation itself is unchanged.

The SMART: Organisational AI Maturity Model

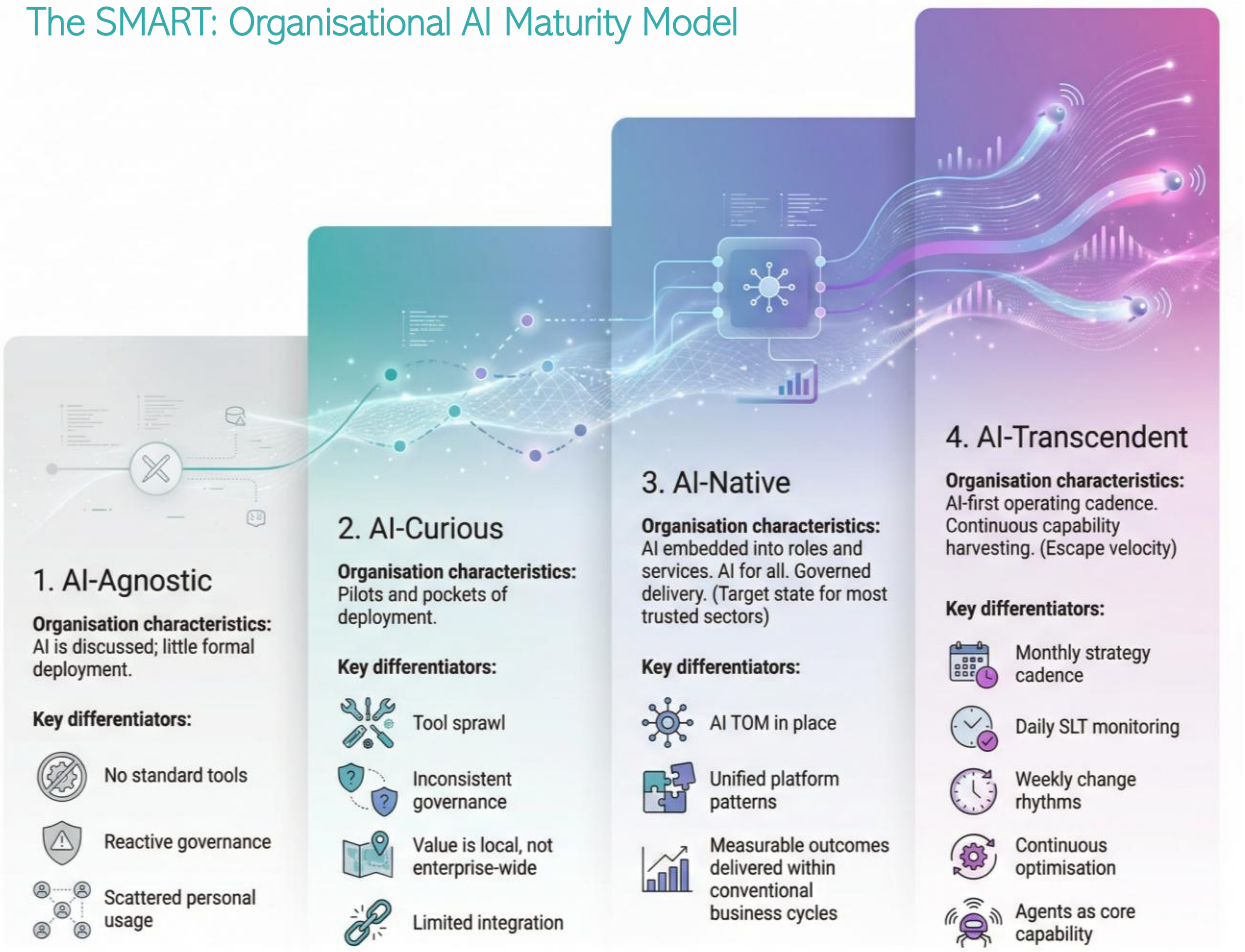

In practice, most trusted and regulated organisations will achieve significant returns at the AI-Native level. Transcendent becomes relevant when competitive or viability pressures demand a step-change in organisational velocity, and when leadership is willing to adopt radically faster decision and change cadences.

This distinction between AI-Native and AI-Transcendent operating modes is explored in more detail in the ICS.AI AI-Native to AI-Transcendent framework [17].

6. WHAT CHANGES AT AI-TRANSCENDENT: THE CADENCE SHIFT

GenAI's most fundamental deliverable is a step-change in individual capability and speed. But organisations only maximise that benefit when their operating cadence allows. If the organisation plans annually, decides slowly, and changes infrequently, it will under-capture weekly improvements in AI capability and workflow design.

AI-Transcendent organisations therefore redesign their business cycles to match the rate of change:

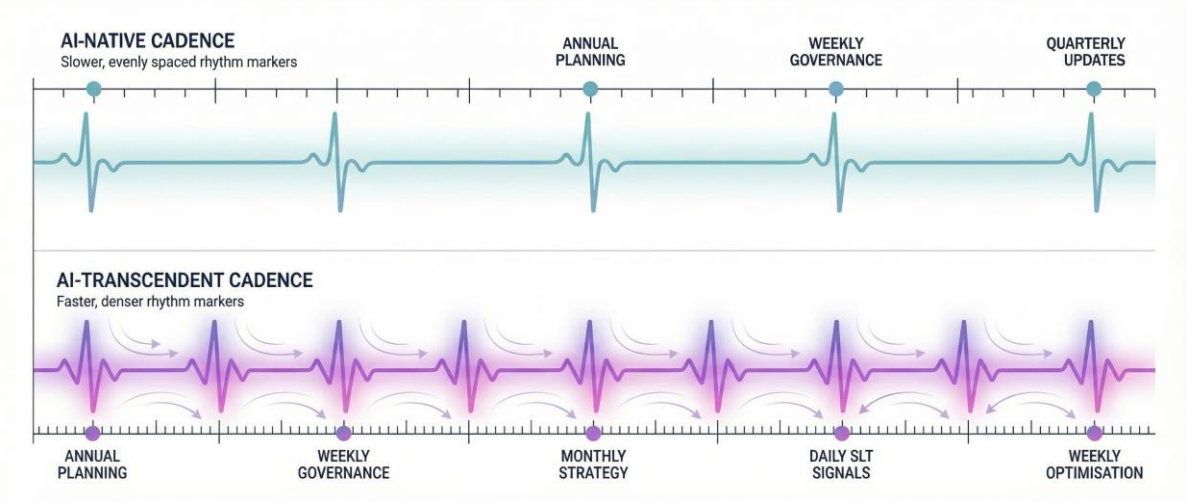

Operating Element	Typical AI-Native Cadance	AI-Transcendent Cadence
Strategy and planning	Annual strategy with quarterly refresh.	Monthly strategy and planning cycles with rapid reprioritisation.
SLT decision-making	Weekly/fortnightly governance forums.	Daily decisioning and monitoring against the value ledger.
Change communication	Monthly updates and milestone comms.	Weekly all-hands change planning.
Skills training	Onboarding plus periodic refresh.	Quarterly in-person depth training and weekly learning loops.
TOM and standards	Updated quarterly or when policy changes.	Updated weekly to reflect tool changes and learnings.
Use case optimisation	Continuous improvement on priority use cases.	Continuous optimisation across portfolio; agents as evolving products.

Transcendent is not 'doing the same work faster'. It is changing the system so that faster capability can compound: more releases, more learning loops, and more improvements captured per month.

7. THE SMART: AI TARGET OPERATING MODEL (AI TOM): THE EXECUTION ENGINE

Whether your ambition is AI-Native or AI-Transcendent, you will not achieve it without an AI Target Operating Model (AI TOM). The AI TOM defines how the organisation will operate with AI as a foundational capability: who decides, how work is done, how knowledge is managed, how risks are controlled, and how value is measured.

The critical maturity step from AI-Curious to AI-Native is the combination of both an AI TOM and an AI for All policy. Without a TOM, AI adoption remains opportunistic. Without AI for All, strategy cannot be executed at scale.

7.1 What the SMART: AI-TOM Defines

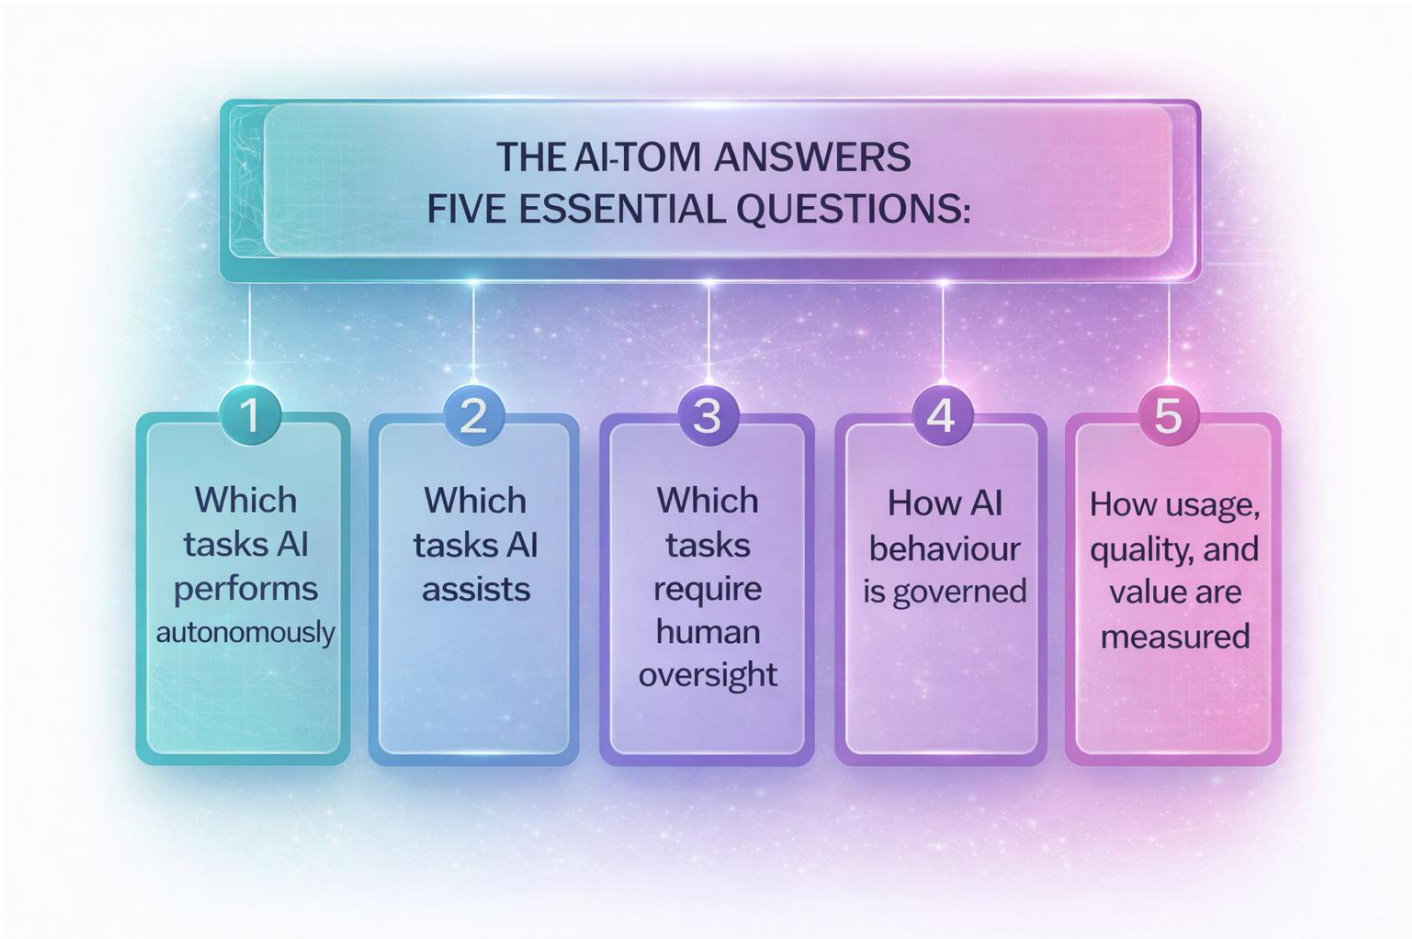

7.2 The Seven Domains of SMART: AI-TOM

A key SMART principle is that the TOM should not remain a slide deck or policy memo. It should be **encoded into the platform itself** through governance guardrails, agent permissions, runtime constraints, escalation rules, and compliance validation. In other words, the operating model becomes executable.

8. THE SMART: UNIFIED AI PLATFORM: ARCHITECTURE FOR SAFE SCALE

The SMART: Unified AI Platform provides the technical foundation that implements the transformation. In architectural terms, it brings together four capabilities:

- AI interaction surfaces for staff and citizens
- Structured organisational knowledge for grounding
- Governance infrastructure for safe AI operation
- Agentic production systems that build and maintain AI services.

The strategic argument for unification is simple: the organisation can only scale AI when the enabling technology reduces friction rather than adding it. Affordability, integration, and compliance-by-default determine whether AI becomes enterprise infrastructure or remains a set of disconnected experiments.

8.1 Platform architecture overview

Layer	Core Components	Strategic Role
Surface Layer	SMART: Front Door (citizen/customer engagement), SMART: Workforce AI (staff productivity), AI-first websites, web chat, phone AI, email, SMS, social, and assisted digital channels.	Provides the interaction surfaces through which citizens and staff engage with AI across every channel.
Knowledge and Governance Layer	SMART: Service Graph and SMART: Digital Dome.	Stores structured organisational knowledge, permissions, guardrails, auditability, compliance controls, and runtime governance.
Production Layer	SMART: Agentic Foundry.	Uses agentic development swarms and reusable patterns to specify, build, manage, monitor, and improve AI services.

8.2 Surface Layer: AI-first engagement

The Surface Layer is where humans interact with AI. It contains two core products: the **SMART: Front Door** for citizen and customer engagement, and **SMART: Workforce AI** for staff productivity.

SMART: Front Door

The Front Door is the AI-driven engagement layer used by citizens or customers. It replaces fragmented engagement channels – CMS-based websites, standalone chatbot tools, contact-centre add-ons, email flows, SMS automation, and social media response tools – with a unified AI engagement layer. All channels – AI-first websites, phone AI, web chat, email, SMS, social, and assisted digital support – are grounded in the same organisational knowledge and governed by the same controls.

AI-first websites represent a new model of digital service delivery. Traditional websites require residents to navigate pages and forms. AI-first websites allow users to describe their need in natural language. The system identifies the relevant service and guides the person through a structured service journey that can include eligibility checks, information capture, validation, escalation, and transaction completion. Traditional browsing remains available, but conversational interaction becomes the primary interface.

SMART: Workforce AI

SMART: Workforce AI provides AI assistance embedded into everyday staff work – drafting, analysis, policy navigation, case summarisation, knowledge support, and process guidance – governed by role-based permissions and organisational policies. Where the Front Door transforms how citizens access services, Workforce AI transforms how staff deliver them.

Together they ensure that AI is not confined to a single channel or user group but operates across the full breadth of the organisation's interactions.

8.3 Knowledge and Governance Layer: structured knowledge plus control

This layer is what separates governed AI from ungoverned AI. It contains two components that together ensure every AI interaction – whether citizen-facing, staff-facing, or agent-executed – operates within defined knowledge and governance boundaries.

SMART: Service Graph

The Service Graph is the structured representation of organisational services, rules, processes, fees, evidence requirements, and legislation. Its purpose is to move knowledge from static documents and web pages into machine-usable form so that AI systems can ground their responses in **authoritative organisational logic** rather than in loosely interpreted fragments or generic training data. The Service Graph is continuously maintained and versioned, ensuring that AI responses reflect current policy and practice.

SMART: Digital Dome

The Digital Dome is the governance control plane. It provides permissions management, escalation rules, audit logging, kill switches, compliance validation, and FinOps monitoring. Every AI action across every Surface Layer channel is subject to the Digital Dome's controls. The governing principle is straightforward: **no autonomy without control**.

In practical terms, this means an organisation can answer the compliance test that trusted sectors require: what did the AI do, why did it do it, what knowledge was it grounded on, and which policy controls were applied – for every channel and every use case.

8.4 Production Layer: the SMART: Agentic Foundry

The Agentic Foundry is the production engine of the platform. It harnesses **agentic development swarms** – coordinated teams of specialised AI agents – to build and maintain AI-first websites, applications, automation agents, and integrations in a governed way suitable for regulated environments. This is where the structural shift described in Section 2 becomes operational: instead of traditional software development timelines measured in months, the Foundry enables continuous creation, improvement, and reconfiguration of digital services.

Its operating lifecycle follows three stages:

1. **Specify:** create a structured Foundry Brief describing services, rules, integrations, and governance constraints. The brief captures what the organisation needs – drawing on the Service Graph for knowledge and the Digital Dome for governance boundaries – in a form that agentic swarms can act on.

2. **Build:** agentic swarms generate and assemble the required AI-first websites, applications, automation agents, and integrations – tested against compliance requirements before deployment.

3. **Manage:** continuously monitor live services, validate compliance, optimise performance, and improve – treating deployed services as evolving products rather than static deliverables.

8.5 The three revolutions SMART: represents

The SMART: platform sits at the intersection of three structural transformations. No single transformation delivers the full prize; it is the combination of all three – enabled by a unified platform – that allows organisations to capture the economic benefits described in this paper

Revolution	What it Means	Platform Component
AI-Native Organisations	Organisations redesigned around AI as core infrastructure – not a bolt-on to existing operating models.	SMART: AI Transformation Programme and AI-TOM encode the target operating model into the platform.
Agentic Software Production	Digital systems built and evolved through AI development infrastructure rather than traditional multi-month software projects.	SMART: Agentic Foundry: Specify → Build → Manage lifecycle using agentic development swarms.
AI-First Engagement	Citizens and customers interact through conversational service journeys rather than navigating traditional websites and forms.	SMART: Front Door: AI-first websites, phone AI, web chat, email, SMS, social, and assisted digital.

8.6 What a unified platform replaces

Most organisations currently operate fragmented digital engagement and automation estates. The SMART: platform replaces these with one unified, governed AI platform:

- Fragmented CMS platforms and separately governed website estates → replaced by AI-first websites built and maintained through the Agentic Foundry.
- Standalone chatbot tools and contact-centre AI add-ons → replaced by the SMART: Front Door with unified knowledge and governance.
- Accessibility and content-governance point solutions → built into the platform's service-delivery workflow rather than operated as separate tools.
- Isolated email and SMS automation tools → unified within Front Door channel management.
- Workflow automation systems requiring bespoke integration → replaced by agentic back-office agents governed through the Digital Dome.
- Multiple disconnected knowledge bases with inconsistent policy interpretation → replaced by the SMART: Service Graph as the single authoritative knowledge source.

8.7 Affordability, integration, and compliance-by-default

Premium AI subscriptions are commonly priced per seat. At scale, even modest per-user pricing becomes a material operating cost. This creates a predictable failure mode: only a small subset of staff receive approved tools, adoption remains thin, and the value ceiling stays low. To execute an organisation-wide portfolio, the unit economics must make universal access viable.

Equally, no trusted organisation will commit to wholesale AI use unless it can prove what the AI did, why it did it, what it was grounded on, and which policy controls were applied – across every channel and use case. If the platform cannot do that, the programme will either stall before scale or proceed without assurance.

9. EVIDENCE FROM PRACTICE: DERBY CITY COUNCIL

Derby City Council is a well-documented UK example of a whole-council AI programme. Derby is a strong pattern match for TOM-led scaling because it exhibits: multi-service deployment (not a single isolated pilot), a portfolio of prioritised opportunities (261 → 54), and clear operational metrics.

Indicator	Reported Figure	Source
Citizen enquiries handled	1.4 million+ questions; ~43% of queries; 24/7	Derby City Council [12]
Deflection rate improvement	20% expected → 43% actual	Computer Weekly [15]
Portfolio approach	261 opportunities identified; 54 prioritised	Computer Weekly [15]
Contract value (expansion)	£7m contract to extend AI across services	Derby City Council [12]
Verified savings to date	£12m identified	Derby Telegraph [16]
Governance	AI Ethics Board (independent chair); DPIAs per use case	Government Technology [14]

Material outcomes appear when the programme is whole organisation, use cases span multiple services, and the enabling platform and governance are designed for reuse and scale.

9.1 Derby AI Transformation – At a Glance [13]

10. ECONOMIC BENEFITS FRAMEWORK: VALIDATED FINANCIAL PROJECTIONS

Moving from case studies to organisation-wide planning requires validated, evidence-based financial projections. This section presents conservative estimates for two categories of economic benefit: **bankable income generation** (revenue that exists but is not fully collected) and **efficiency savings** (reduced costs through productivity and automation). All figures use lower-bound assumptions validated against independent research, UK government trials, and the Ministry of Housing, Communities & Local Government Revenue Outturn 2024-25 data [24].

10.1 Bankable Income Generation

AI-enabled revenue maximisation targets income that already exists but is not being fully collected or optimised. This is not about increasing tax rates or cutting services - it is about ensuring all liable parties are correctly identified and billed, optimising collection, and identifying opportunities within policy constraints.

Council Type	Conservative Low	Conservative High	% of Net Revenue
District	£0.63m	£1.40m	1.9%
Unitary	£3.49m	£8.01m	1.2%
County	£6.88m	£15.32m	0.9%

Primary income levers: Council tax discount verification and arrears prediction; business rates hereditament discovery and relief verification; adult social care financial assessment automation; HRA arrears prevention and void reduction; planning validation and CIL/S106 optimisation; parking demand forecasting; treasury cashflow optimisation.

10.2 Efficiency Savings

Efficiency savings require careful consideration of how productivity gains are converted into cashable budget reductions - typically through vacancy management, contract renegotiation, or redeployment of capacity.

Council Type	Validated Low	Validated High	% of Net Revenue
District	£0.92m	£1.38m	5.6% - 8.5%
Unitary	£4.43m	£10.86m	1.7% - 4.1%
County	£8.89m	£24.13m	1.1% - 3.1%

10.3 Key Validation: Workforce AI Savings

The commonly cited claim of 'one day per week' savings (7-8 hours) for staff productivity tools is higher than independent evidence supports. The UK Government trial of 20,000 civil servants found 2.2 hours per week savings (26 minutes/day) [10][18]. Similar results from Somerset Council and Forrester supports a validated range of 2-4 hours per week . [19][20]

Independent academic evidence also shows that GenAI assistance can increase worker productivity by around 14–15% in real work settings, reinforcing the plausibility of these ranges [26]. Use this range for business case planning; treat higher figures as stretch targets.

10.4 Strongly Validated Use Cases

The following use cases have robust, independent evidence demonstrating material impact when deployed at scale:

- Contact Centre AI: McKinsey analysis indicates 30-45% cost-reduction potential in customer operations through AI-enabled automation, triage, and assisted resolution [21]. UK local government evidence aligns with this range. Contact centre benchmarks show that automation and AI-enabled self-service can handle over 40% of inbound interactions for suitable call types, which council case studies report material FTE-equivalent capacity release from AI-driven contact systems [23][25].

- Social Worker Administrative Support – The British Association of Social Workers (BASW) reports that up to 65% of social worker time is spent on administrative tasks rather than direct practice [22]. Independent evaluations of ambient voice and AI documentation tools show reductions in documentation time of approximately 20-30%, with several minutes saved per client interaction. At scale, these gains translate into meaningful workforce capacity recovery [27][28][29].

- AI-Enabled Citizen Contact System – Socitm-validated case studies, including Hillingdon Council, demonstrate that AI-driven citizen contact platforms can deliver sustained demand deflection, improved first-contact resolution, and measurable cost savings, where deployed as part of whole-system service model rather than isolated pilots [25].

- Contact Centre Benchmarking and Performance Evidence – UK contact centre benchmarks confirm that organisations combining AI self-service, assisted handling, and intelligent routing consistently outperform peers on cost per contact, handling time, and service accessibility, providing a reliable baseline for benefits modelling [23].

10.3 Implementation Approach

A phased 24-month approach delivers 60% of benefits within the first 12 months:

- **Phase 1 (Months 0-6):** Foundation - core platform, citizen chatbot, staff copilot pilot, priority revenue agents. Expected benefit: 20-30%.
- **Phase 2 (Months 6-12):** Revenue Protection - full council tax and business rates suite, HRA collection agents. Cumulative benefit: 60%.
- **Phase 3 (Months 12-18):** Service Income - planning automation, parking optimisation, adult social care charging. Cumulative benefit: 85%.
- **Phase 4 (Months 18-24):** Optimisation - grant intelligence, treasury, cross-domain orchestration. Cumulative benefit: 100%.

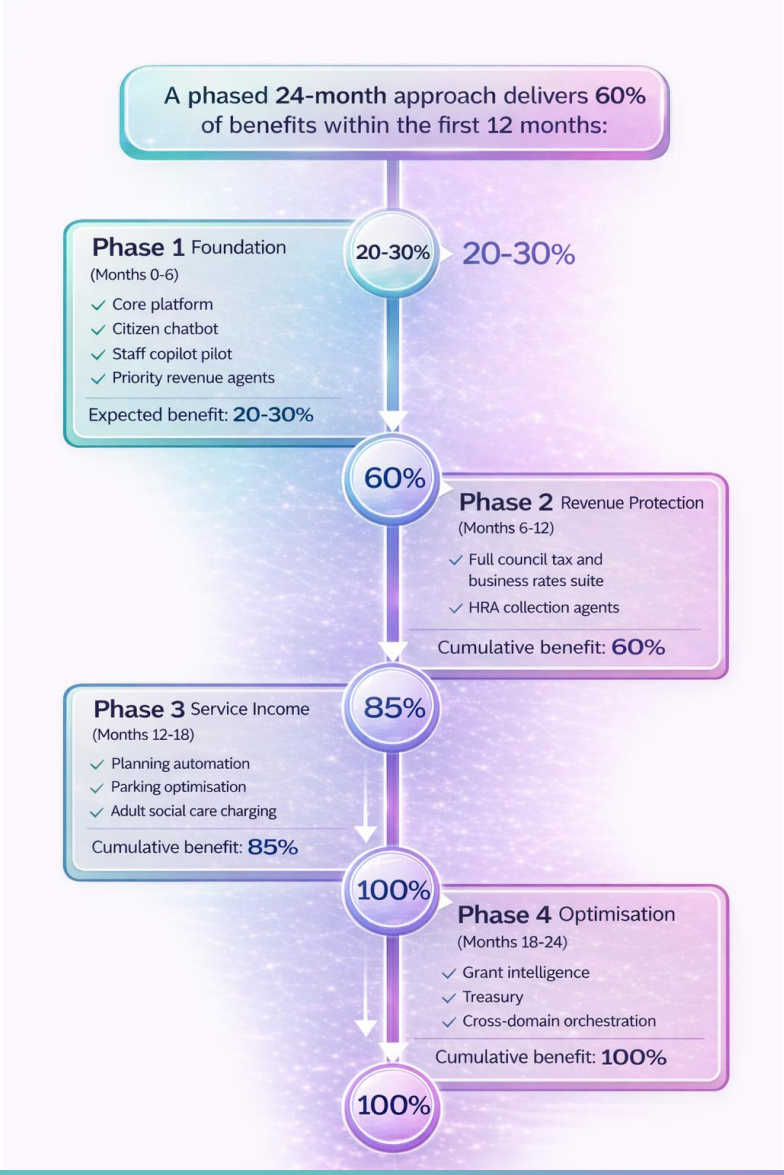

11. MEASURING SUSTAINABLE VALUE: THE AI VALUE LEDGER

To avoid 'pilot theatre'*, value measurement must be designed in from day one. A practical mechanism is an **AI value ledger**: a transparent list of AI-enabled changes with baselines, owners, expected impact, real-world performance, and benefit realisation status. Remember: in the LGA survey, 56% of councils 'don't know' whether they've achieved cost savings from AI [9]. The ledger solves this.

"Pilot theatre" refers to the repeated running of small AI pilots that generate activity and learning but are not designed or governed in a way that allows value to scale, be measured, or be sustained.

11.1 Value Categories

Category	Typical Metrics	Examples
Demand reduction	Deflection %, first-contact resolution, handling time.	Enquiry resolution by AI; reduced call waiting.
Workforce productivity	Time saved per role, throughput, cycle time.	Report drafting; case summarisation; policy navigation.
Process automation	End-to-end cycle time, cost per transaction, error rates.	Debt recovery; document processing; triage.
Quality and outcomes	Decision consistency, compliance rates, satisfaction.	Consistent advice; faster responses; improved experience.
Income protection	Collection rates, fraud detected, income recovered.	Single Person Discount (SPD) verification; arrears prediction; property discovery.

In a mature AI programme, the value ledger is not a reporting afterthought. It is the mechanism that links the maturity target, the TOM, the platform, and the portfolio of use cases to financial and service outcomes that leadership can govern.

12. ROADMAPS

12.1 Minimum Viable Path to AI-Native

1. **SLT sets the ambition**
   (Native) and measurable outcomes required.
2. **Define and approve the AI TOM**
   Decision rights, governance, policies, training, measurement.
3. **Select an affordable platform**
   That enables AI for All across all three domains.
4. **Roll out AI for All**
   With mandatory training, instrument usage and quality.
5. **Deliver first wave of high-volume use cases**
   Prove measurement via the value ledger.
6. **Scale the portfolio**
   Using reusable patterns; embed continuous improvement.

12.2 Path to AI-Transcendent

1. **Commit to cadence change**
   Monthly strategic cycles, weekly change rhythm, daily SLT monitoring.
2. **Treat standards as living products**
   Weekly updates to grounding, prompts, rules, agent behaviours.
3. **Invest in quarterly depth training**
   Build internal 'AI coaches'.
4. **Expand agentic automation**
   Standardise integration patterns.
5. **Run continuous optimisation**
   Measure, refine, redeploy; capture compounding gains.

13. RISKS AND CONTROLS FOR TRUSTED SECTORS

Moving fast increases both the upside and the obligation to control risk:

- **Data protection:** role-based access, data classification, leakage controls, audit logging.
- **Quality and safety:** grounding in authoritative content, clear escalation, human-in-the-loop.
- **Bias and fairness:** monitor outcomes, test for disparate impacts, update policies.
- **Model drift:** controlled change and testing process for AI updates.
- **Workforce risk:** transparent comms, training, and role redesign.
- **Vendor lock-in:** design for model flexibility and portability.

14. PLATFORM EVALUATION CHECKLIST

Platform Evaluation Checklist

1	Can we afford to provide approved AI to every staff member without rationing?
2	Does it span all three domains (front door, copilot, agentic) with common governance?
3	Can we reuse components across services for fast deployment?
4	Does it provide auditable logs, policy enforcement, and role-based controls?
5	How does it handle grounding and content governance?
6	Does it support multi-model routing?
7	What is the operating burden: new systems, admins, contracts?
8	Can we measure value end-to-end?

14. CONCLUSION

AI can deliver transformative value, but only when it is scaled safely and deliberately. The evidence is unambiguous:

- 78-88% of organisations are now using AI [2][3]
- Yet only 31% are scaling enterprise-wide [3]
- 74% have yet to demonstrate tangible value [4]
- The gap is not technology – it is operating model

Local government will not reach material outcomes through isolated pilots alone. Councils need a senior-leadership decision on maturity target, a clearly defined AI Target Operating Model to translate ambition into execution, and an affordable unified platform that provides AI for All across the Front Door, Workforce AI, and the agentic back office.

AI is becoming part of the operating system of the organisation. The SMART: platform embodies three structural revolutions

1. AI-Native Organisations redesigned around AI
2. Agentic Software Production that replaces traditional development timelines
3. AI-First Engagement that transforms how citizens access services.

The Service Graph structures organisational knowledge for machine use; the Digital Dome ensures no autonomy without control; and the Agentic Foundry builds and maintains digital services at a pace that was previously impossible. In trusted sectors, that future only works when governance is embedded directly into the digital infrastructure.

The economic prize is substantial. For a typical unitary authority, the combined opportunity from income generation and efficiency savings is **£8m - £19m per year** [24] with a larger range for county councils. Derby City Council demonstrates that a whole-organisation, governed approach can already produce measurable savings [12][16].

In an environment where AI use is already happening, the real choice is not whether to allow AI – but whether to provide a sanctioned, governed, and value-driven way to use it.

The path is proven.

What remains is the leadership decision to begin.

15. REFERENCES

[1] NAO. Local government financial sustainability. Feb 2025.

[2] Stanford HAI. AI Index Report 2025.

[3] McKinsey. The State of AI: Global Survey 2025.

[4] BCG. AI Adoption in 2024: 74% of Companies Struggle to Achieve Value. Oct 2024

[5] Gartner Predicts 30% of Generative AI Projects Will Be Abandoned by End of 2025. Gartner Newsroom. 2024.

[6] Bain. Executive Survey – AI Moves from Pilots to Production. 2024.

[7] IBM. Global AI Adoption Index 2023.

[8] NAO. Use of AI in Government: Key Facts. March 2024.

[9] LGA. State of the Sector: AI. 2024 and 2025 Update.

[10] GOV.UK – Department for Science, Innovation and Technology – Microsoft 365 Copilot Experiment Cross-Government Findings Report.

[11] GOV.UK. NHS Microsoft 365 Copilot trial.

[12] Derby City Council. Portfolio Spotlight. Dec 2024.

[13] ICS.AI Derby City Council Case Study. Jan 2026.

[14] Government Technology. Derby City Council story.

[15] Computer Weekly – Derby City Council eyes £12.25m annual savings with AI. July 2024

[16] Derby Telegraph AI saves Derby City Council £12 million and fills 100 jobs. Jan 2026

[17] ICS.AI. AI-Native to AI-Transcendent: Operating cadence and maturity model. December 2025

[18] UK Government Digital Service (GDS). M365 Copilot Experiment Report. June 2025.

[19] Forrester Consulting. Total Economic Impact of Microsoft 365 Copilot. March 2025.

[20] LGA Case Studies: Somerset Council. Feb 2025

[21] McKinsey. The economic potential of generative AI. 2023.

[22] British Association of Social Workers(BASW) - 80-20 Campaign Report.

[23] ContactBabel. UK Contact Centre Decision-Maker's Guide. 2024.

[24] Ministry of Housing, Communities & Local Government - Local authority revenue expenditure data. 2024-25.

[25] Socitm. AI-driven citizen contact system - Hillingdon Council Case Study – January 2025.

[26] Brynjolfsson, Li & Raymond. Generative AI at Work. NBER/QJE - November 2023.

[27] Topaz, M. et al. *Navigating the uncharted risks and benefits of AI scribes in clinical practice.* Digital Medicine, 2025.

[28] Stults, C.D. et al. *Evaluation of an ambient artificial intelligence documentation platform in clinical settings.* JAMA Network Open, 2025.

[29] Great Ormond Street Hospital NHS Foundation Trust. *Ambient Artificial Intelligence Phase 4 Evaluation Report.* NHS England, 2025.

ABOUT THE AUTHOR

Martin Neale is CEO of ICS.AI, the UK's leading AI transformation company for local government.

ICS.AI has achieved significant market share in council AI solutions, with deployments across more than 20 UK local authorities generating demonstrated savings in multiple councils.

The company's SMART: platform and AI Target Operating Model (AI-TOM) methodology enable organisations to move from 'AI-Curious' to 'AI-Native' status through a structured transformation framework that delivers guaranteed savings while preserving and redeploying workforce capability.

CONTACT

WEBSITE

https://www.ics.ai/

ICS.AI LINKEDIN

https://www.linkedin.com/company/ics-ai-ltd

MARTIN NEALE LINKEDIN

https://www.linkedin.com/in/martinneale/

TAKE A LOOK AT OUR OTHER WHITE PAPERS

The Human Firewall: Why AI won't replace most jobs

Why the 'End of Work' is a Myth - and Why the Real Challenge is Accountability, Not Obsolescence.

ICS.AI