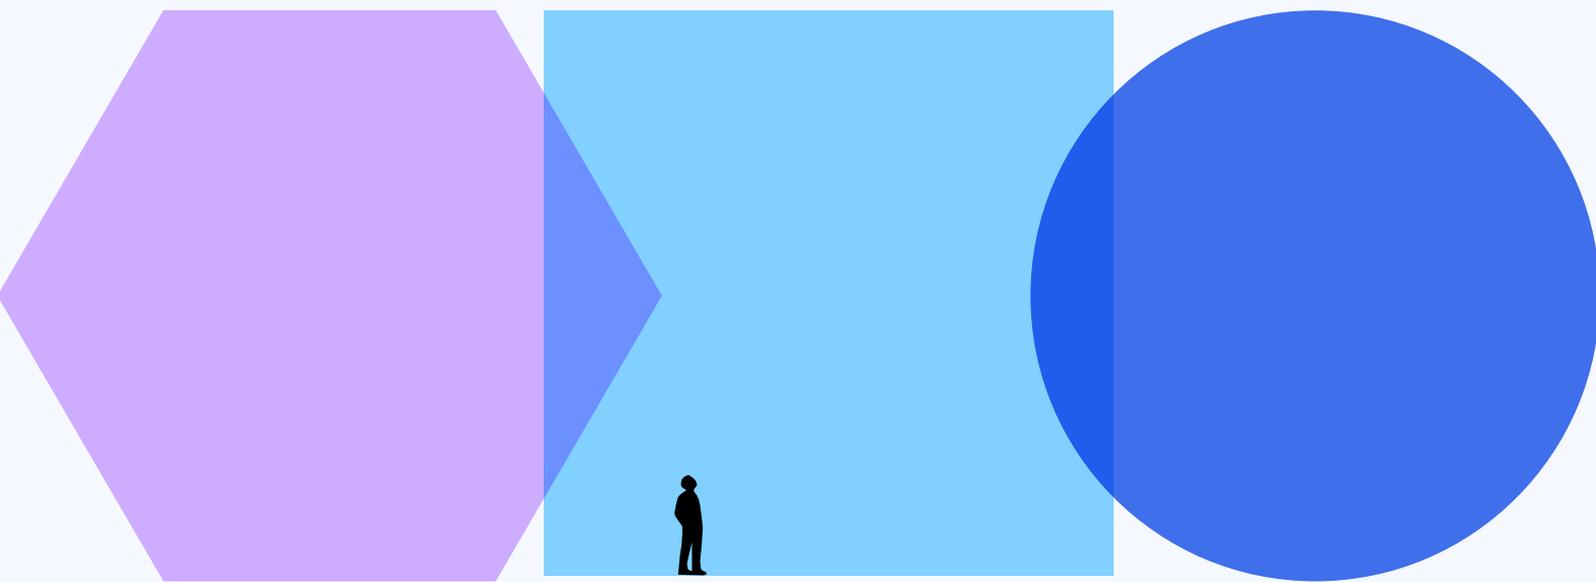


Elsewhen Reports

The Agentic State: A Blueprint for Public Sector Productivity

Nadav Mordechai & Gabriel O'Brien



Contents

Executive Summary	3
Our Research	4
5 Key Themes	5
Redefine public sector productivity	6
Centre experimentation on real problems and real value	8
Rethink the technical foundations for Agentic AI	11
Adopt pragmatic approaches to governance and workforce debate	14
Learn from regulated industries, not just the public sector	16
The Agentic Blueprint for Public Sector Productivity	18
Phase 1 – Map	19
Phase 2 – Build	21
Phase 3 – Scale	23
Phase 4 – Evolve	25
Working with Elsewhen	27

Executive Summary

UK public sector productivity has been stagnant for three decades. Costs have risen. Demand has intensified. Performance has yet to return to pre-pandemic levels. Productivity is now the central test of fiscal sustainability and service reform.

Artificial Intelligence is increasingly presented as the solution. The economic case is substantial. The Office for Budget Responsibility estimates that effective AI adoption could unlock up to £41 billion annually. The Tony Blair Institute suggests returns of roughly £9 for every £1 invested. The opportunity is significant. But capturing it will require structural change, not incremental efficiency.

To date, adoption has focused on assistive tools: drafting copilots, document summarisation and task automation. These generate measurable time savings. But public sector productivity is defined by outcomes delivered per pound spent. Faster briefings do not reduce backlogs by default. More efficient administration does not automatically improve case resolution. Without redesigning services, system performance remains largely unchanged.

Evidence increasingly points to a different approach: agentic AI. These are autonomous, goal-oriented systems capable of coordinating workflows across data, tools and teams. Rather than improving isolated tasks, they operate across services. The shift is from individual efficiency to institutional performance.

This report sets out what that shift requires in practice. Drawing on primary interviews and wider research, we introduce a four-phase Agentic Blueprint for Public Sector Productivity – a structured approach to deploying agentic AI with clarity, accountability and measurable impact. The core question is not whether AI can save time. It is whether it can materially improve how public services deliver value.



Nadav Mordechai,
Director of Product & Strategy,
Elsewhen

Our Research

This paper is grounded in original research combining secondary analysis of published policy, economic data, and academic literature with in-depth interviews with experts from academia, government, consultancy, and policy, including:



Karl Hoods,
Chief Digital
and Information
Officer at ICS Digital



Dr Nina Jorden,
Assistant Professor
at the Bennett School
of Public Policy, Univer-
sity of Cambridge



Vahid Panjgani,
Head of Data & AI at
Elsewhen



Niall Grant,
Public Sector
Transformation
Lead at Clarasys

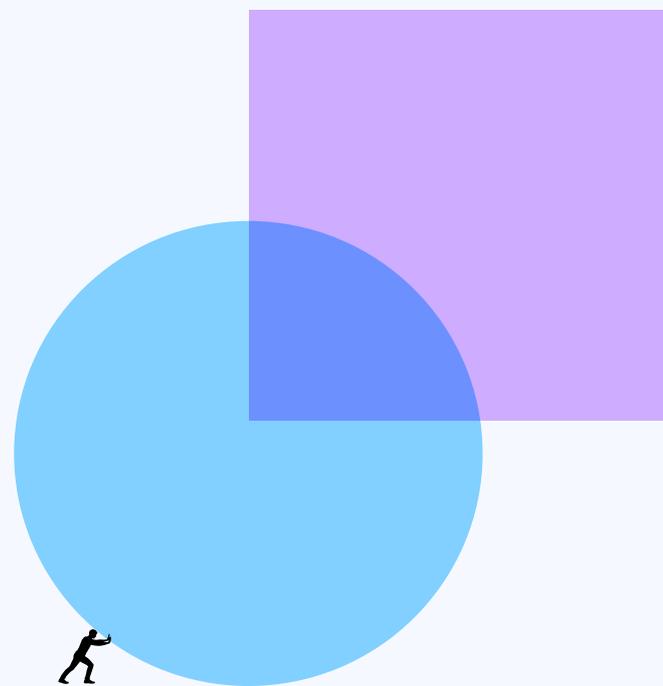


Kat Excell,
Senior Product
Manager at the
Tony Blair Institute

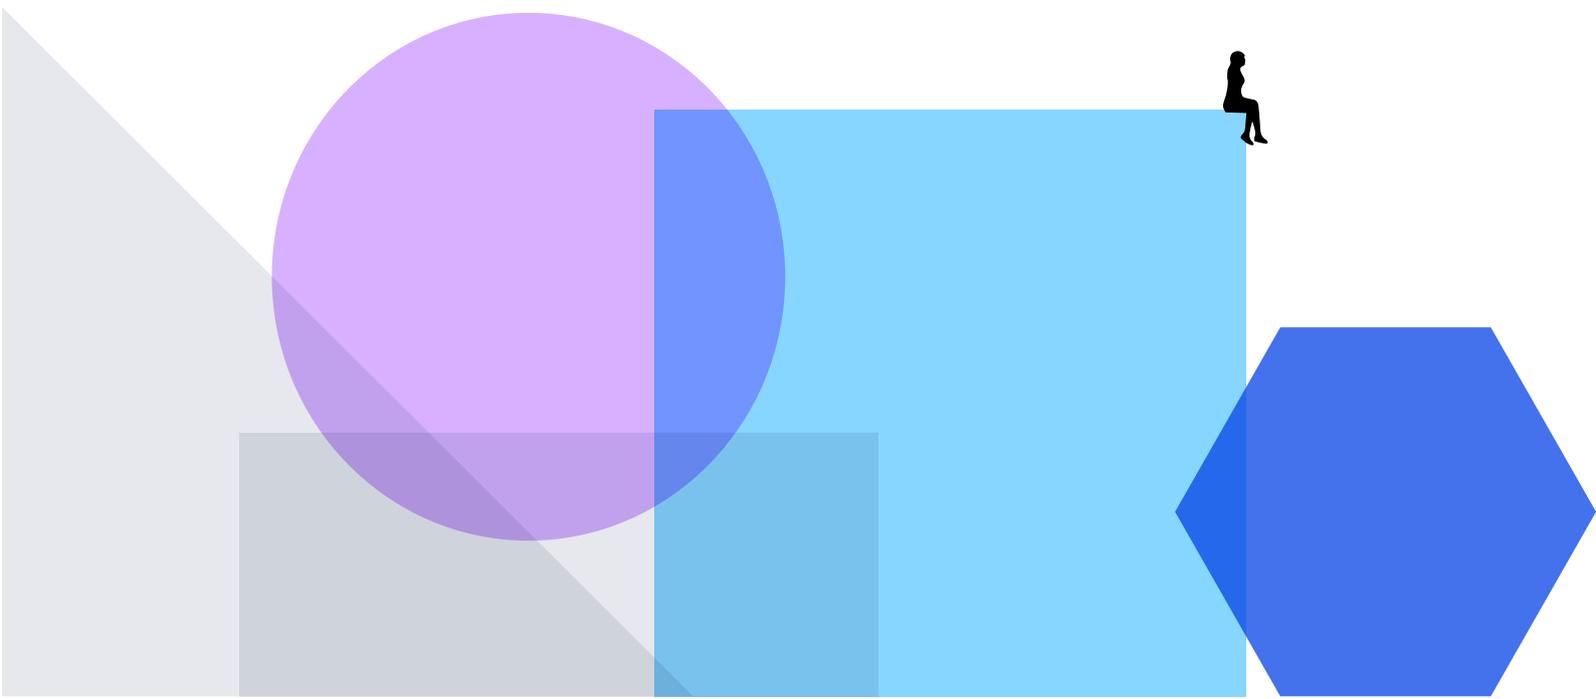
Throughout our conversations, a clear consensus emerged: agentic AI demands more than incremental adoption. It requires a rethinking of how government measures value, designs services, governs decisions, and develops its workforce. We want to thank all of the participants for their time contributing to this work.



Gabriel O'Brien,
Researcher at Elsewhen

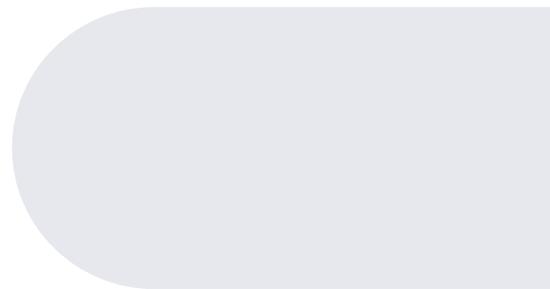
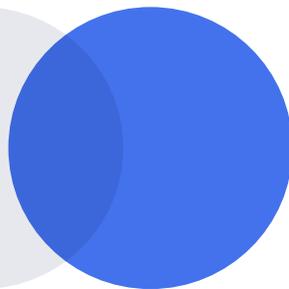
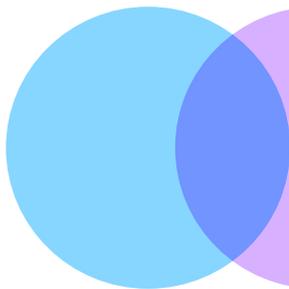


5 Key Themes



01

Redefine public sector productivity



Our research revealed a fundamental disconnect between how public sector productivity is measured and what actually matters. Current metrics, designed for counting activity, including forms processed, cases closed, hours logged, tell leaders nothing about whether citizens are better off. This creates a cycle where the wrong problems are diagnosed and the wrong solutions applied.



For just over 40% of public services in the UK, official statistics still assume 'output = input', which automatically means zero productivity growth. Even when better measurement methods are available, they often fail to take into account preventive measures, inter-agency dependencies or quality improvements." — Dr Nina Jorden, University of Cambridge

Interviewees consistently supported this, highlighting how flawed measurement drives flawed decision-making. When metrics reward volume, organisations optimise for volume, adding headcount or cutting costs rather than questioning whether the work itself delivers value.



We are measuring success in the wrong way, focusing on output rather than outcome. 'Problems' are framed as issues of too many people or we are not moving fast enough. And then, projects tend to focus on incremental change around the edges rather than fundamentally reimagining services to most effectively meet policy outcomes." — Niall Grant, Clarasys

In an AI-first world, this redefinition becomes even more urgent. When machines own the activities we currently count, the entire basis of measurement collapses. Without new frameworks, there is no way to assess whether AI is actually improving anything that matters.

We see a strong case for an alternative view of productivity: a whole-system view that captures not just outputs but service flow, citizen outcomes, and workforce capacity.

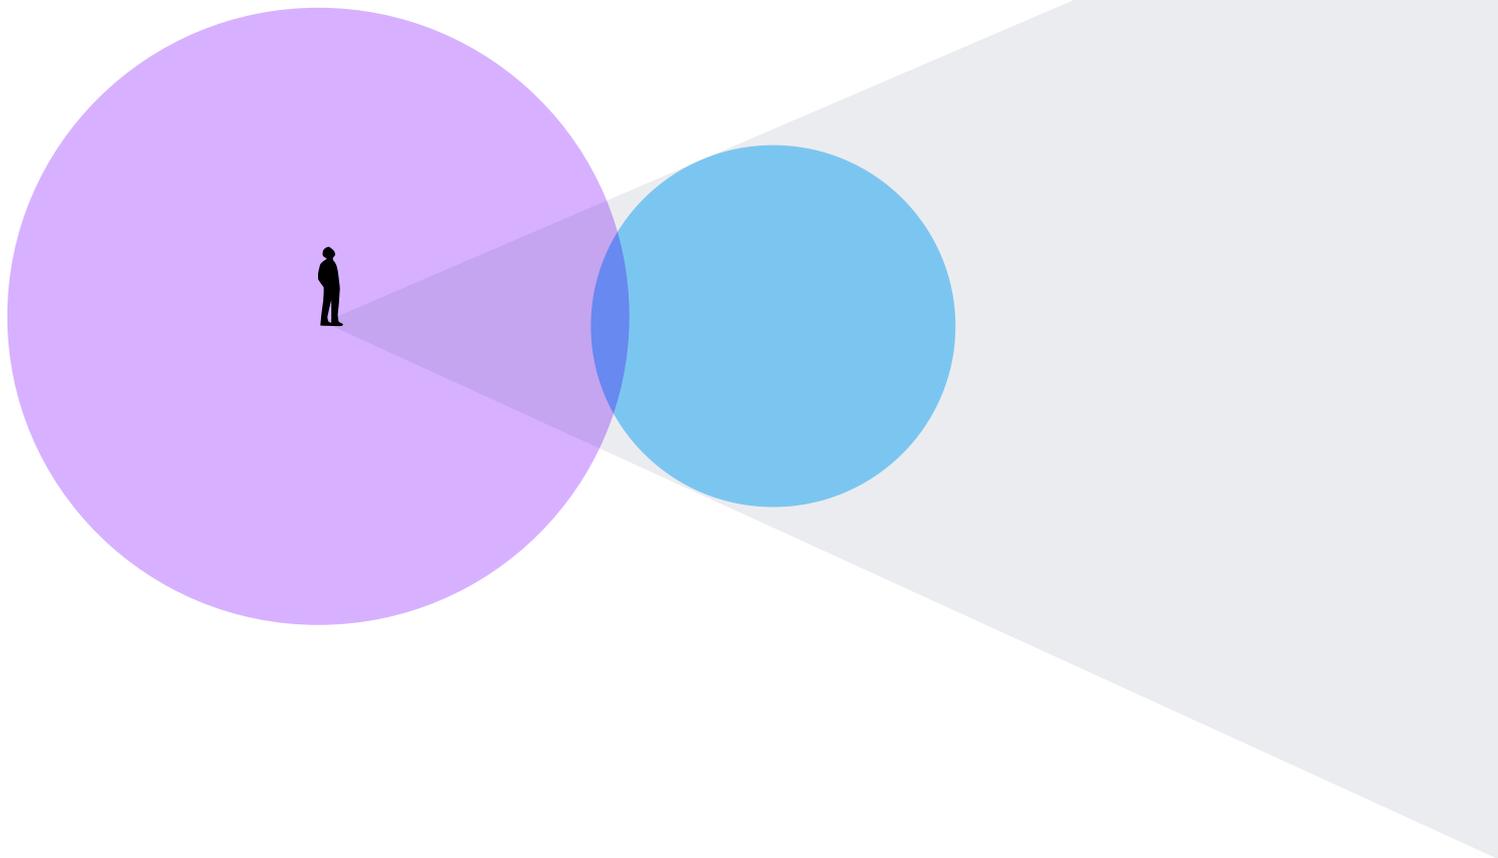


A future productivity framework should treat productivity as a whole-system outcome, not a departmental metric. AI enables much more integrated operations, so measurement should reflect that. A modern framework would assess not just outputs, but system flow, service experience, quality and equity of outcomes, and the capacity and wellbeing of the workforce. This moves productivity from a narrow cost/output focus to one that recognises accuracy, human capability, and how effectively demand moves through the system."

— Kat Excell, Tony Blair Institute

02

Centre experimentation on real problems and real value



When experimenting with agentic AI, the focus must be on the problem, not the technology. The question should not be "where can we deploy AI?" but "what are we trying to achieve, and what intervention would actually help?"



We're thinking less about the technology and what the latest development might be at the moment and really getting people to focus on the problem statement, which brings in two other tracks. One is that we're trying to get our wider audience to start thinking from a systems perspective and the second is to promote a product mindset and really be focused on what we are trying to achieve. With this we can determine whether or not agentic AI, machine learning, automation, whatever it may be, will help deliver that outcome."

— Karl Hoods, ICS Digital

Whilst focusing on the problem, our interviews also pointed towards clear patterns where agents should be mobilised to add the most value; specifically, those tasks that consume significant staff capacity while requiring little genuine judgement.



Agentic AI has the greatest potential where civil servants face repetitive, manual, and time-consuming tasks: searching for information across systems, updating case notes, checking eligibility rules, compiling templated reports, or analysing routine datasets. These are high-volume, low-judgement activities that consume significant capacity." — Kat Excell, Tony Blair Institute

Identifying AI opportunities is only useful, however, if organisations can prioritise them systematically and separate what's strategically worthwhile from what's merely technically possible.



We've developed what we call the TRIAL framework. It's a systematic way to evaluate whether a workflow is a good candidate for agentic AI. What's useful about this framework is it forces you to look beyond just 'can we automate this?' and instead ask 'should we automate this, and what would success look like?' It helps separate the genuinely high-value opportunities from things that might be technically possible but not strategically worthwhile."

— Vahid Panjganj, Elsewhen

But identifying the right problem and the right opportunity is not enough. Our research highlighted that across the public sector, experiments routinely demonstrate possibility without delivering operational impact. The problem starts with how success is defined.



I've got a problem with the term proof of concept because anyone can prove or disprove any concept, but how do we know it's adding value? So we are actively now starting to shift to think about how we approach these experiments from a Proof of Value perspective." — Karl Hoods, ICS Digital

Even well-intentioned pilots fail to scale when they're built in isolation from the operational reality they're supposed to improve.



Are these PoC teams working on real data? Is it fresh data? Do they have the blessing of the higher-ups? Are they working with real scenarios, real situations, real customer data? Or is it separate from reality? If they do PoCs separate from reality because they want to be independent and flexible, then the value that's proved has a big delta to impact because in a real scenario, it wouldn't work the same way."

— Vahid Panjgani, Elsewhen

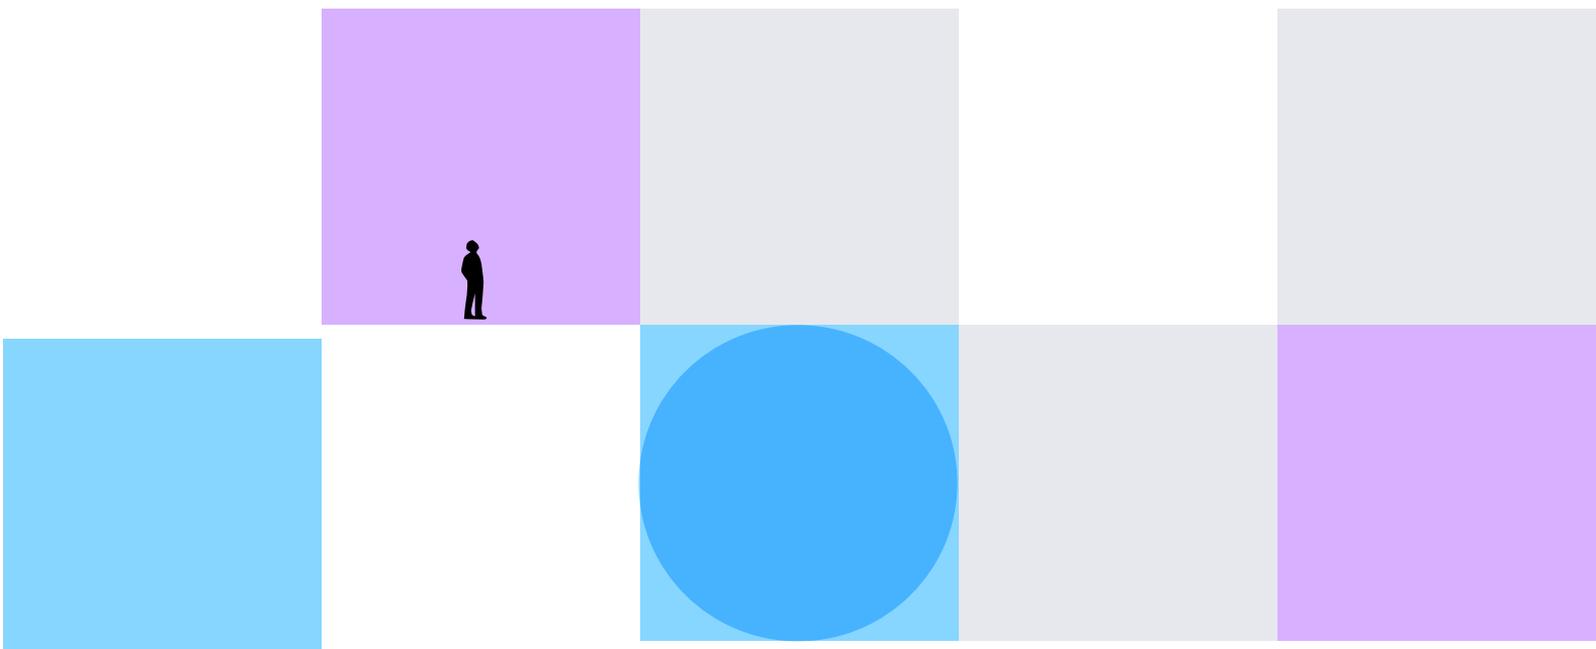
This is a system-wide issue, we see this phenomenon across the public sector. Ultimately, without connecting experiments to genuine delivery problems, innovation activity substitutes for actual reform.



The pilot trap exists because AI is often treated as an innovation exercise rather than genuine delivery reform. Ministries run isolated experiments to show progress, but these rarely connect to real operational problems or to each other." — Kat Excell, Tony Blair Institute

03

Rethink the technical foundations for agentic AI



The emergence of Agentic AI is now challenging industry norms on data maturity and legacy technology. For instance, on legacy technology, what emerged from interviews was a clear message: legacy and AI adoption can, and must, run in tandem.



The 'legacy tech issue' is too big, and will take too long to solve, to be a barrier - transformations are going to have to be run in tandem. We're deploying AI on legacy tech while also tackling the 'legacy tech issue'. So it's actually more of an opportunity than two things in contrast. If the onus and drive is there to make it happen, these things can happen quickly — there just needs to be acceptance of dual running"

— Niall Grant, Clarasys

What's more, the nature of legacy itself has shifted. Even in more digitally mature organisations, cloud migration hasn't removed the legacy problem, it has created a new version in the form of platform sprawl across hyperscalers rather than ageing on-premise systems.



Legacy is one of those problems that can be poorly defined. For example, we did most of our move to the cloud eight or so years ago. So ours is a new legacy problem, we've now got a bunch of other services that we need to keep updating. The solution should be a mindful choice of tooling, governance that goes around it and then continuous improvement. This is increasingly important in this world of agents and AI, where the market is shifting and changing so quickly."

— Karl Hoods, ICS Digital

A potential enabler of this more mindful approach has emerged: an orchestration layer that sits across vendors, giving organisations visibility and control as they deploy agents across increasingly complex platform landscapes.



If we're consuming AI from multiple vendors in our ecosystem, how do we genuinely understand what's happening behind the black box of that? And is there a separate approach where we might actually think we want some kind of middleware orchestration layer or something which we've got more insight and control over?"

— Karl Hoods, ICS Digital

Increasingly, the conventional wisdom that organisations must fix their data before adopting AI is being challenged. Evidence is actually emerging that agents often need something different entirely.



I often hear leaders say 'we don't have data readiness, we don't have great data'. But I don't think agents need data in the formal shape that was crucial for traditional AI. So we are actually dealing with a new breed of data gap." — Vahid Panjganj, Elsewhen

Instead, what agents need is the kind of knowledge that has never been formalised, the information people carry that enables them to actually do their jobs.



There's the data that's in front of you on your screen — formal, structured, registered in a system. Then there's the knowledge that lives in people's heads: which team to escalate to, how a particular minister likes to be briefed, when to push a decision and when to hold back, which stakeholders need to be in the room. Everyone has their own version of it, it's what actually drives how work gets done and is exactly what agents need access to. This data together is what enables effective decision-making."

— Vahid Panjganj, Elsewhen

This new breed of data shifts our focus from prompt engineering to context engineering, where organisations structure and expand the knowledge and business context that agents draw on.



It becomes more around context engineering than prompt engineering. How do we build that knowledge context within our domain to be able to access it with agents?"

— Karl Hoods, ICS Digital

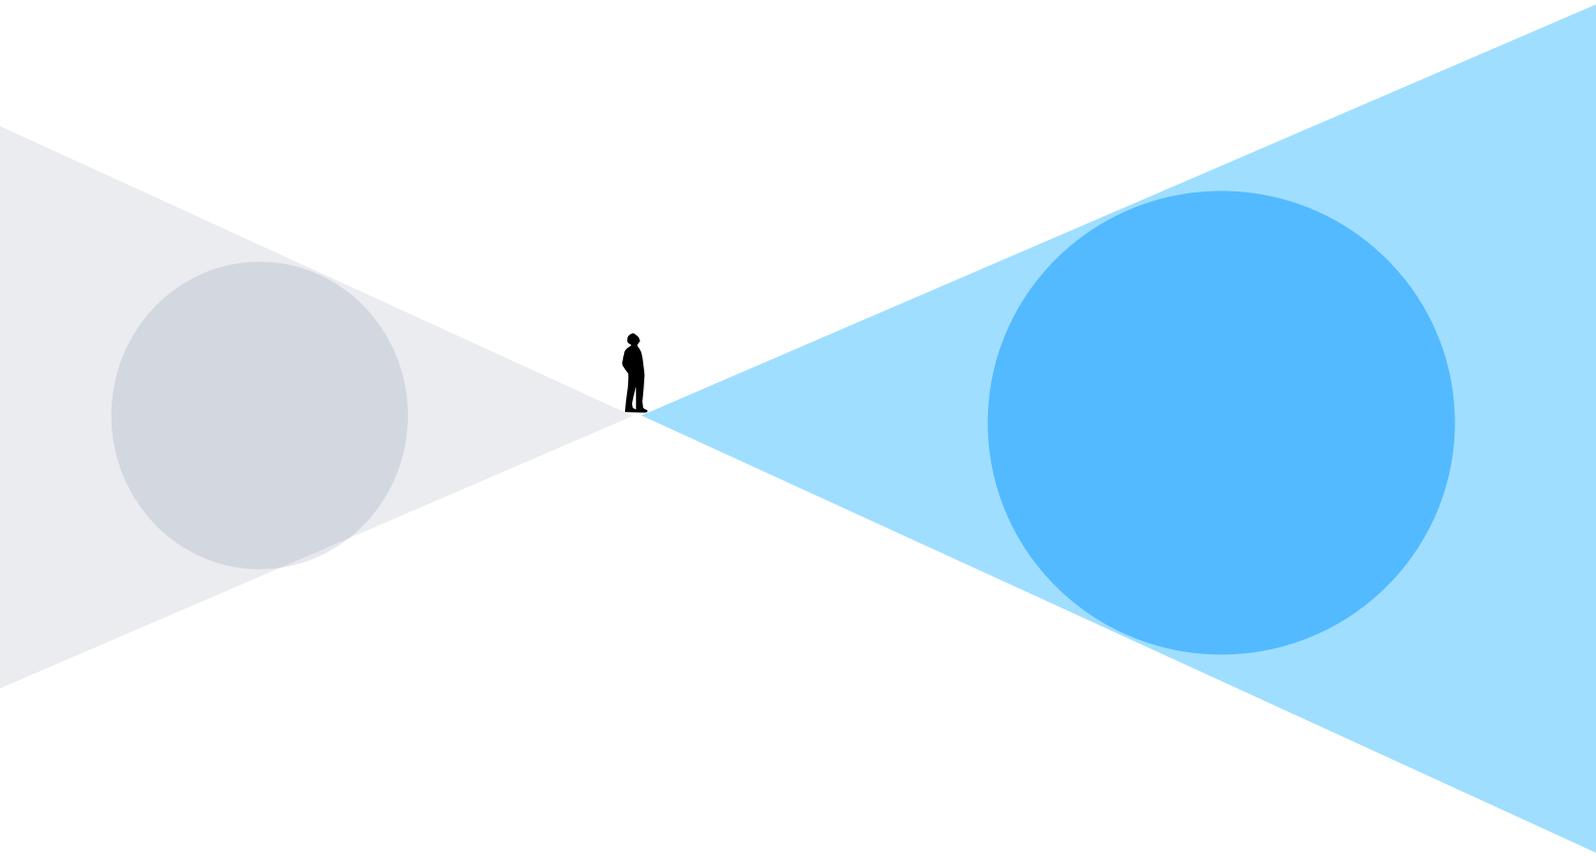
The research also surfaced emerging technical approaches to this challenge – specifically, modelling relationships that traditional databases were never designed to hold.



Graph databases are emerging as a critical enabler for agentic AI because they model relationships between people, policy, tasks, precedents and best practice to map who owns what, which policy governs which decision, what sequence of actions led to a successful outcome. Traditional relational databases were never designed to hold that kind of interconnected knowledge. This takes agents from simple task execution to complex, context-aware decision-making across an entire service chain." — Vahid Panjgani, Elsewhen

04

Adopt pragmatic approaches to governance and workforce debate



Governance and workforce impacts are, rightly, prominent concerns around agentic AI in the public sector. But, often, the theoretical debates risk blurring and overcomplicating a practical path forward.

For example, given the existing governance frameworks in the public sector, the instinct to build entirely new governance structures for AI is largely unnecessary. Privacy assessments, security protocols, data handling procedures are sufficient. However, the governance gaps are perhaps narrower and more specific than most organisations assume.



We've actually got a lot of the governance we need around this sort of activity. Processes around privacy impact assessments, privacy notices, where data's stored, the security that wraps around it, we have the same processes regardless of what the technology is that you're using. The bit that we're mindful of is how are we assuring ourselves that we've got insight into the decisions that are being made? So if you're using AI to make a set of decisions, do you understand what the algorithm is doing, what data sources it took, and then how it arrived at that decision? That's the bit of difference that we're looking at at the moment and deploying, rather than creating a separate AI governance workstream."

— Karl Hoods

Similarly, on workforce evolution, we can see practical approaches such as analysing tasks within roles, not roles themselves, and planning for hybrid teams where agents handle specific structured activities alongside human colleagues.



"The focus really does need to be on how they evolve roles, not just replace them as we risk losing crucial deep contextual knowledge from the workforce."

— Niall Grant, Clarasys



We are actively trying to think about tasks within job descriptions and roles where we think AI and agentic AI has a role to play. Using this, we can eventually embed that into workforce planning so you almost end up with this hybrid human-machine type workforce planning perspective."

— Karl Hoods

In fact, the evolution of roles through AI can actually reduce administrative burden and cognitive overload and create a more stable and more experienced.

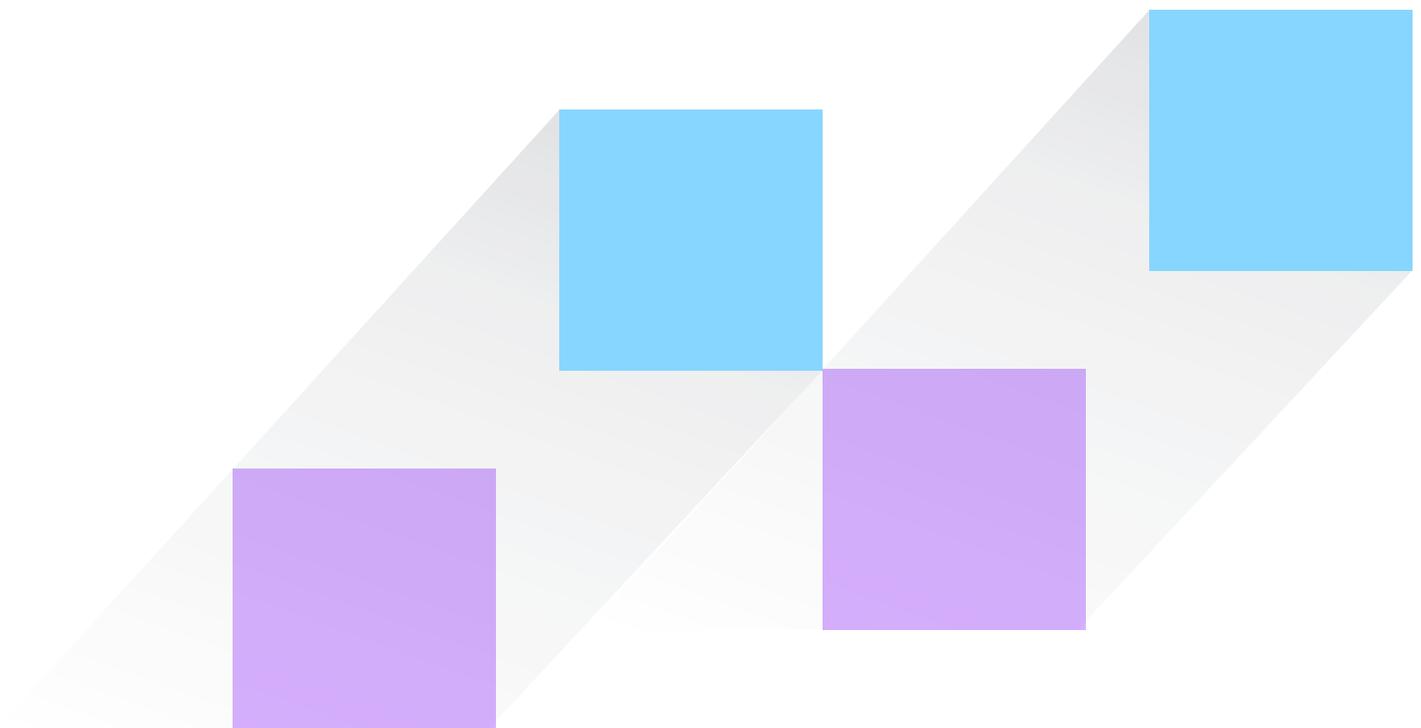


AI can reduce administrative burden and cognitive load and if that genuinely translates into lower turnover and fewer absences, the same labour costs can buy more experienced, effective working hours and more embedded organisational know-how. But that outcome isn't automatic. With agentic AI, some teams are seeing the opposite: work expands, pace accelerates, and people can burn out as they try to 'manage' tireless digital co-workers. Capturing the benefits without the costs requires deliberate work design, such as protecting breaks and reflection time, and ensuring the work isn't going too fast."

— Dr Nina Jorden

05

Learn from regulated industries, not just the public sector



The public sector has a tendency to draw inspiration from other public bodies but ignore highly regulated private sector organisations including banks, insurers, energy companies, telecommunications providers. These organisations, who face comparable demands around compliance, data governance, algorithmic accountability, and public trust, have often moved further and faster in deploying AI at scale than the public sector.



Having worked across multiple sectors you can get fixated on the sector you're in and all the constraints you have, and I hear things like 'We don't want to talk to a bank because they've got billions to invest.' But they're also highly regulated. So if they're able to make a difference in that regulated environment, how can we learn from that and take a percentage of what they do and apply it to our work?" – Karl Hoods

This isn't just a theoretical argument, cross-sector learning should be a regular, deliberate practice rather than an occasional exercise.



The majority of challenges that the public and private sector face are pretty similar, just the context in which we're operating in is the difference. In ICS Digital, on average once a month, the leadership team visits an organisation – we don't mind what sector they're in – for example, we've had everything from Formula One teams through to banks and insurance companies. to create connections and share knowledge." – Karl Hoods

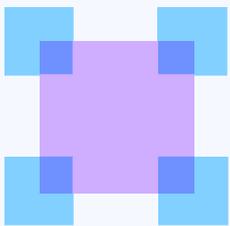
Taken together, these themes lead to a clear conclusion. Agentic AI is not just a technology rollout. It requires changes to how public sector organisations are designed and run.

Productivity will not improve through isolated pilots or marginal efficiency gains. It will improve when organisations measure the right outcomes, link experimentation to delivery, modernise their technical foundations, strengthen practical governance, and redesign roles around hybrid human-machine teams.

The Agentic Blueprint for Public Sector Productivity

The Agentic Blueprint for Public Sector Productivity is built from Elsewhen's experience delivering AI and digital transformation alongside public sector organisations. It translates what our research uncovered into a sequenced delivery path, four phases, each designed to build the conditions for the next.

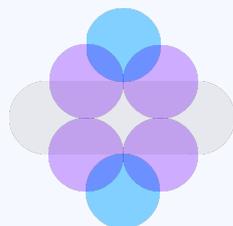
The four phases — Map, Build, Scale and Evolve — are not a maturity model. They are a practical sequence for adopting agentic AI with clarity, control, and impact. Each phase includes specific actions grounded in what we've seen work in practice.



Map

Lay the foundations before deploying technology

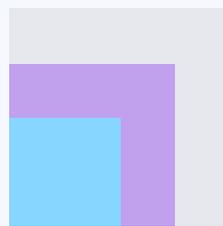
- Redefine Productivity
- Identify High-Value Opportunities



Build

Ground AI in delivery reality and real data.

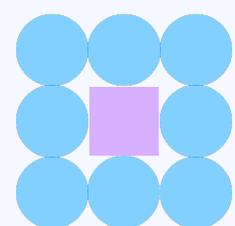
- Prove Value, Not Concept
- Build Orchestration Across Platforms



Scale

Move from pilots to institutional performance.

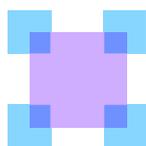
- Optimise Governance for Explainability
- Scale with Confidence



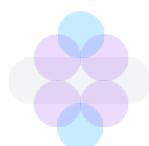
Evolve

Redesign services and processes around agents.

- Build Organisational Memory
- Establish a Hybrid Operating Model



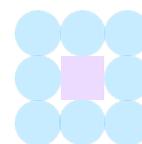
Map



Build



Scale



Evolve

Phase 1 – Map

Understand where you are, where the real opportunities lie, and establish the foundations for agentic AI.

Most organisations skip these foundations. The pressure to demonstrate AI progress to ministers, boards and the public incentivises visible deployments over foundational work. The result is predictable: scattered portfolios of pilots that prove possibility without delivering impact, AI tools that optimise tasks nobody questioned in the first place, and efficiency gains that look good in a briefing but don't move the needle on outcomes that matter to citizens.

The MAP phase exists to prevent this. It establishes the foundations that every subsequent phase depends on: a redefined view of productivity that captures outcomes, not just outputs, and a systematically prioritised pipeline of opportunities where agentic AI can deliver genuine transformation.

From our experience, we suggest the following activities:

Redefine Public Sector Productivity

Redefine productivity from activity metrics to outcome-based measures that capture prevention, quality, citizen wellbeing, and systemic value creation.

For public sector leaders wondering what this could look like in practice, our interviewee, Dr Nina Jorden from University of Cambridge outlined a four step process for building a productivity framework fit for the AI era:

- Map the service chain — Every service provider should be able to map their service chain with key figures ranging from budgets to inputs to outputs and outcomes. This allows AI projects to be assessed in all areas.
- Adjust outputs for quality and welfare — Output indices should, where possible, be adjusted for quality and welfare. The contribution of AI is then assessed on the basis of its impact on these adjusted outputs and not just on the basis of pure volumes.
- Track organisational learning and agility — As management and human resource practices are important drivers of productivity, the framework should track indicators of organisational learning and agility, such as investment in training, digital skills, collaborative practices and the presence of an agile workforce capable of adopting and adapting technologies. These are the capabilities that enable AI to move from experimental pilot projects to scaled, system-wide use.
- Add metrics for responsible use — A framework for the AI era needs metrics for responsible use: transparency, handling of bias, data quality, and trust from employees and users.

Identify High-Value Opportunities

Use evidence-based frameworks to identify high-leverage workflows where AI delivers genuine transformation rather than marginal efficiency gains.

Deploy structured evaluation frameworks combining qualitative insight with quantitative signals to score each opportunity systematically. Elsewhen's TRIAL framework assesses candidates across five dimensions:

Tedious

Form-filling, reconciling, logging, scheduling. Activities that drag on people's time and energy. Free people from this drag and keep humans for judgment and relationships.

Repetitive

Patterned, high-frequency work. If it's boringly predictable, it's probably agentic.

Inefficient

Seemingly "fine but slow" processes that compound across teams. Agents can turn hours into minutes without cutting corners.

Accessible

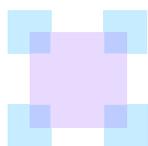
Tasks where systems and data are reachable. If an agent can't access the information, it can't act.

Low-risk

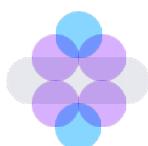
Begin where consequences are limited. Prove value quickly, then scale into adjacent, higher-value flows.

You can use this framework in our AI Opportunities Workshop [available here](#).





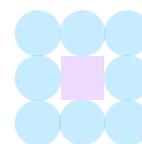
Map



Build



Scale



Evolve

Phase 2 – Build

Move from analysis to action. Build the evidence and technical foundations for enterprise-wide agentic AI.

In the public sector, pilots are often designed as innovation exercises, disconnected from the operational problems they're supposed to solve. They run on sanitised data, measure the wrong things, and stall before delivery. Meanwhile, technical strategies default to one of two extremes: wait until legacy systems are modernised before deploying AI, or layer AI on top of fragmented platforms with no visibility into what's happening underneath. Neither works.

The BUILD phase addresses both. It grounds experimentation in operational reality and creates the architectural foundations that allow agentic AI to operate across complex, multi-vendor environments.

From our experience, we recommend the following activities:

Ground AI Pilots in Value and Reality

Structure pilots on real data with clear paths to scale, measuring impact through leading indicators tied to policy outcomes and citizen value.

At Elsewhen, we are proactively trying to solve the 'pilot trap' for organisations. Rather than running isolated experiments that prove a concept but stall before delivery, we help organisations build the skills, resources and space for dedicated AI Squads.

These cross-functional teams are embedded directly into departments, focused on proving value from Agentic AI with real data, in real teams, on mission-critical problems.

These squads work alongside operational staff from day one, ensuring that what gets built reflects the actual complexity of the service rather than a sanitised lab version of it. The goal is to compress the distance between experiment and production: identify a high-value workflow, prove its impact in weeks, and create a clear path to scale so that pilots become the first deployment, not a dead end

Build Orchestration Layers to Unlock Existing Technology and Data

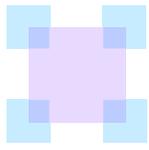
Navigate legacy technology by strategically managing existing platform capabilities while building orchestration layers that enable agentic systems.

Legacy technology and imperfect data are not reasons to delay. But deploying agents across a landscape of Microsoft, AWS, Salesforce, ServiceNow and others — each with their own AI capabilities — without visibility or control is a recipe for platform sprawl, duplicated effort, and vendor lock-in.

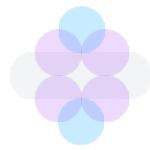
At Elsewhen, we are helping organisations solve this problem by building an orchestration layer that sits across existing platforms and gives organisations visibility and control as they deploy agents. This agentic architecture is made up of three components:

- A shared knowledge layer — extracting only the facts, relationships, and context that agents need to make decisions.
- An agentic sourcing layer — querying data where it lives, extracting knowledge from documents and systems without forcing migration.
- A consumption layer — exposing this knowledge consistently to the people, tools, and agents that need it.

The result is the ability to deploy agents across complex platform landscapes with genuine insight into what's happening — without locking into any single vendor or waiting for perfect modernisation.



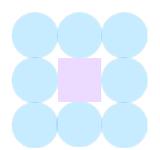
Map



Build



Scale



Evolve

Phase 3 – Scale

Move beyond individual deployments towards an agentic enterprise model, guided by governance that enables rather than blocks.

The transition from successful pilots to enterprisewide deployment is where most public sector AI programmes stall. Not because the technology fails, but because the organisational conditions for scaling were never established. Individual teams prove value in isolation, but without shared patterns, coordinated governance, and a deliberate approach to cross-service adoption, each new deployment starts from scratch, duplicating effort, fragmenting investment, and reinforcing the silos that agentic AI is supposed to break down.

The SCALE phase is not about doing more pilots faster. It is about creating the organisational infrastructure that allows proven capabilities to spread, with control, consistency, and confidence.

From our experience, we recommend the following activities:

Optimise AI Governance Around Explainability and Human Oversight

Enhance existing governance frameworks with AI-specific considerations for algorithmic explainability and human oversight, rather than building a separate AI bureaucracy.

When it comes to Agentic AI, the Public Sector will be expected to promote values of fairness, legitimacy, public trust and transparency. Similarly, there is a growing expectation the AI outputs are screened by skilled, human-in-the-loop oversight. However, recent research from LSE highlights a key challenge: many existing frameworks for explainability and human-in-the-loop oversight, developed by frontier AI companies, do not inherently promote fairness, legitimacy or public trust. They are often implemented as generic safeguards rather than tailored to the specific accountability needs of public services.

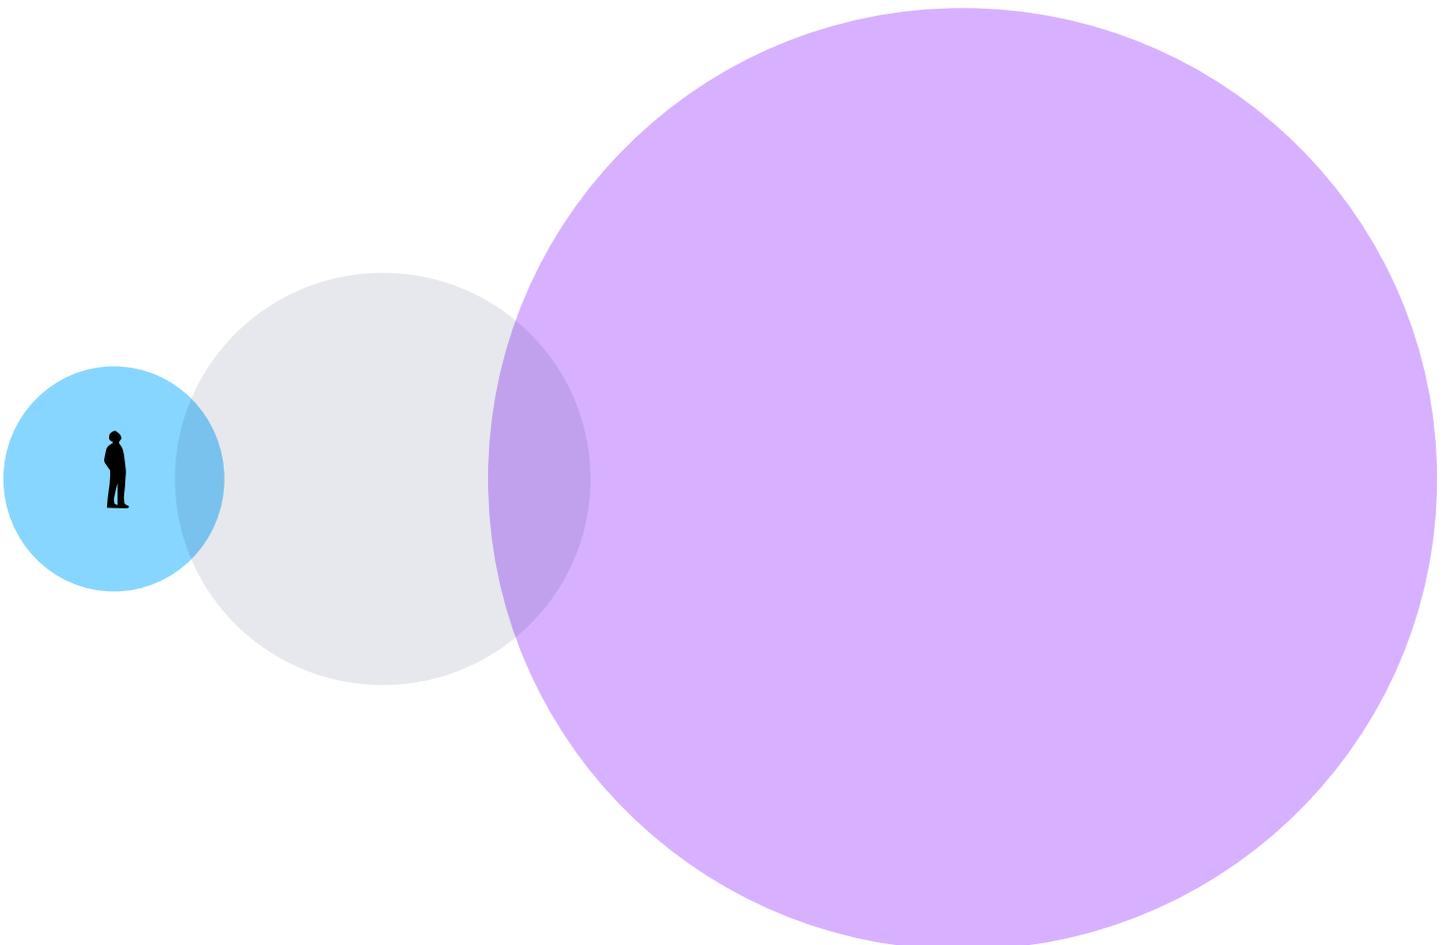
For public sector leaders, the advice is don't default to audit trails or human oversight for their own sake. Ask what each explanation needs to achieve; is it enabling a citizen to challenge a decision, helping a caseworker trust a recommendation, or satisfying a regulator? Design explainability around that purpose, not around a generic checklist.

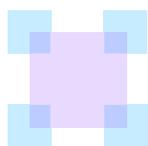
Learn from Regulated Industries to Scale with Confidence

Draw on proven approaches from highly regulated private sector organisations who have already navigated the governance, data, and trust challenges the public sector now faces.

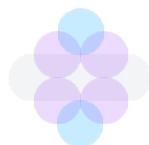
One of the barriers to confident scaling is the belief that the public sector's constraints are unique. In practice, highly regulated industries — banking, insurance, energy, telecommunications — operate under comparable demands around data governance, compliance, algorithmic accountability, and public trust, and have moved further and faster in deploying AI at enterprise scale. Risk management frameworks, governance models for explainability, approaches to human oversight in automated workflows — these have been developed and stress-tested under constraints that closely mirror what the public sector faces.

In practice, this means establishing regular knowledge exchange with these industries — visiting banks who have deployed AI in compliance-heavy workflows, learning from insurers who have tackled algorithmic explainability, partnering with telecom providers who have redesigned roles around AI agents. At Elsewhen, we support organisations to connect across industries with our event series, Elsewhen Rebase. To get involved, get in touch with [Yafaa Ahres](#).

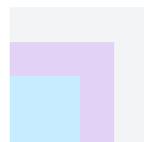




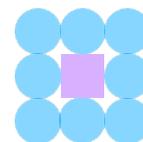
Map



Build



Scale



Evolve

Phase 4 – Evolve

Redesign services and operating models around agentic AI from the ground up, unlocking entirely new forms of value.

This is where the blueprint moves from optimising what exists to reinventing how services work. Not automating current processes, but redesigning them end-to-end around what agents and humans each do best. It is also where organisations confront a deeper data challenge, not cleaning up what they have, but building what they've never had.

When it comes to Agentic AI, the Public Sector will be expected to promote values of fairness, legitimacy, public trust and transparency. Similarly, there is a growing expectation the AI outputs are screened by skilled, human-in-the-loop oversight. However, recent research from LSE highlights a key challenge: many existing frameworks for explainability and human-in-the-loop oversight, developed by frontier AI companies, do not inherently promote fairness, legitimacy or public trust. They are often implemented as generic safeguards rather than tailored to the specific accountability needs of public services.

For public sector leaders, the advice is don't default to audit trails or human oversight for their own sake. Ask what each explanation needs to achieve; is it enabling a citizen to challenge a decision, helping a caseworker trust a recommendation, or satisfying a regulator? Design explainability around that purpose, not around a generic checklist.

Learn from Regulated Industries to Scale with Confidence

Draw on proven approaches from highly regulated private sector organisations who have already navigated the governance, data, and trust challenges the public sector now faces.

Most organisations are not here yet. But the decisions made in Map, Build, and Scale determine whether they can get here. Without outcome-focused measurement, there's no way to assess whether redesigned services are actually better. Without proven deployments and orchestration, there's nothing to redesign around. Without scalable governance, there's no confidence to let agents operate at the heart of service delivery.

From our experience, we recommend the following activities:

Explore New Breeds of Data: context engineering and organisational memory

Invest in context engineering and organisational memory systems that capture informal knowledge alongside traditional data quality improvements.

Organisations need to create a shared knowledge layer that captures business logic, decision context, and the relationships between people, policies, tasks, and practices in a form that agents can reason over.

At Elsewhen, we are helping organisations move beyond traditional databases towards graph-based models where facts, summaries, and entity relationships are extracted from source systems and documents and represented in a way that can be queried, importantly without forcing data to move from where it already lives.

Crucially, this includes formalising the informal: the decision rationale, the communication norms, the institutional knowledge that currently exists only in people's heads. When this knowledge is modelled and accessible, agents can do more than execute tasks, they can make contextual decisions, identify dependencies, and surface insights that no single system could reveal on its own.

Establish a clear hybrid, human-machine operating model

Design for workforce evolution by mapping tasks within roles that can be augmented, treating agents as team members rather than replacement tools.

At Elsewhen, we call this hybrid, human-machine operating model the agentic enterprise. It involves complete process and service reinvention, not automating what exists, but redesigning workflows end-to-end around what agents and humans each do best. This might include reinventing roles and job descriptions with agents in mind: mapping tasks within roles to identify where agents can take on structured, repetitive activities and free people for judgement, relationships, and creativity.

It could involve redesigning processes, where agents handle eligibility checks, case routing, data reconciliation, and progress chasing. And it means embedding this into workforce planning so that capacity models reflect hybrid teams, not just headcount, as well as investing in complementary capabilities, such as training, digital skills, collaborative practices, that enable people to work effectively alongside agents.

Working with Elsewhen

This paper has set out what our research tells us about the structural, cultural, and technical foundations that determine whether agentic AI delivers genuine transformation in the public sector.

Elsewhen works with government departments and public sector organisations to put these ideas into practice. We help organisations navigate each phase of the Agentic Blueprint – from redefining productivity and identifying high-value opportunities, through building orchestration layers and ground-ing pilots in operational reality, to establishing the data foundations and hybrid operating models that make agentic AI sustainable at scale.

If you are exploring any of these challenges, we would welcome a conversation.

Get in touch:
+44 (0)20 3758 3224
hello@elsewhen.com
elsewhen.com



