



Public Sector AI Perspectives: The realities of AI adoption and implementation

For what comes next

tlt.com



Contents

Foreword	3
Executive summary	4
1. The current state of public sector AI	6
2. Governance needs to move at the speed of delivery	8
3. Procuring AI is becoming the first line of AI control	10
4. Trust depends on defensibility, not messaging	12
5. Resilience is part of the licence to operate	14
6. Workforce capability will determine who delivers real value	16
7. The next wave: autonomy and exposure	18
Conclusion	19
How TLT can support	20

Foreword

Artificial intelligence has quickly moved from a future ambition to a present-day priority for public sector organisations.

Across government departments, local authorities, regulators and blue-light services, leaders are exploring how AI can help deliver better services, improve productivity, support decision-making and meet growing demand in increasingly challenging circumstances. Yet while enthusiasm for AI is high, the path to meaningful adoption remains far from straightforward.

Through our work with public sector organisations, we see a common challenge emerging. The question is no longer whether AI has potential, but how organisations can deploy it in a way that is effective, accountable and trusted. Success depends on much more than selecting the right technology. It requires strong governance, clear ownership, resilient operating models, informed procurement and a workforce equipped to work alongside AI.

To better understand how organisations are approaching this challenge, we have engaged with leaders across the public sector and spoken directly with clients who are already applying AI in practice.

This report explores the realities of public sector AI adoption and identifies the practical steps organisations can take to move from experimentation to execution.

The organisations that create lasting value will be those that build the legal, operational and governance foundations needed to scale with confidence. AI is not simply another technology programme. It represents a fundamental shift in how public services are designed, delivered and governed. The decisions made today will help define the future of public service delivery for what comes next.



The question is no longer whether AI has potential, but how organisations can deploy it in a way that is effective, accountable and trusted.”



Gareth Oldale

Partner and Head of Digital,
Data and Commercial, TLT

Executive summary

The findings in this report are drawn from research with 100 public sector leaders across local government, central government, regulators and blue-light organisations, alongside interviews with clients actively applying AI within their organisations.

Together, these perspectives provide a detailed picture of how the public sector is progressing from AI experimentation towards operational deployment.

The research reveals a sector with significant ambition but varying levels of readiness. Public sector leaders recognise AI's potential to improve efficiency, enhance decision-making and deliver better outcomes for citizens and service users. However, many organisations continue to face barriers that extend beyond the technology itself. Legacy systems, fragmented data, skills shortages,

governance complexity and uncertainty around legal and regulatory obligations are all influencing the pace of adoption.

Three themes emerge consistently throughout the findings. First, the challenge is increasingly organisational rather than technical. The most successful deployments are being supported by clear accountability, strong leadership and effective governance.

Second, trust and defensibility are becoming critical factors in determining whether AI can be scaled. Organisations must be able to explain decisions, demonstrate appropriate oversight and maintain confidence among stakeholders, regulators and the public.

Third, long-term success will depend on workforce capability, resilience and the ability to embed AI into day-to-day operations rather than treating it as a standalone innovation initiative.



The public sector now faces a strategic choice: treat AI as a series of pilots at the edge of the organisation or use it to fundamentally modernise how government and public services function. Those that create lasting value will be the ones that align AI to mission-critical priorities, strengthen governance early and build the legal and operational foundations for scaled adoption.”




Bill Hull

Partner and Head of
Government and Public
Services, TLT


What the research shows

Three findings across all sectors that will shape the future of public sector AI.




Ambition is high

Over half of public sector leaders expect AI to improve efficiency and productivity.



Confidence remains low

Nearly half of leaders lack confidence that they can meet legal and regulatory requirements for AI deployment.



Transformation is coming

Two-thirds expect AI to significantly reshape workforce roles and skills.

1. The current state of public sector AI

Public sector organisations are not short of ambition.

By 2030, they expect AI to be contributing to improved efficiency and productivity (67%), better decision-making through data and insights (62%), reduced costs or better use of resources (60%), faster service delivery (56%) and improved citizen or service-user experience (53%).

While public sector organisations are clear on the potential value of AI, many are struggling to put the foundations in place to realise it. Skills shortages remain the most cited barrier to adoption (45%), but governance is emerging as a critical challenge and a competitive advantage for those who get it right. While 34% of organisations lack the governance and assurance frameworks needed to support AI deployment, and 56% report governance-related delays, the most mature organisations are treating governance not as a control mechanism but as an enabler of faster, safer and more scalable AI adoption. The underlying challenge is fragmentation.

More than half (52%) of respondents say systems do not integrate effectively across teams or departments, while almost one in five (18%) report no integration at all. Siloed platforms are cited by 53% as a barrier to collaboration and efficiency, and 51% point to a lack of shared data.

The impact is clear: duplication of effort remains common across the public sector. Nearly three in five respondents (58%) say work is duplicated often or very often, suggesting fragmented systems and disconnected data continue to drive inefficiency and prevent organisations from realising the full benefits of AI.

5 main barriers to AI Adoption:

1. Lack of skills and/or capacity.
2. Legacy systems.
3. Lack of governance and assurance frameworks.
4. Regulatory or legal uncertainty.
5. Data quality/integration issues.



TLT Takeaway

The research reveals a clear shift in the public sector AI conversation. The challenge is becoming less about the technology itself and more about governance, accountability and operational readiness.

Organisations can no longer treat governance as a compliance exercise. With 34% citing a lack of governance frameworks as a barrier to adoption and over half pointing to governance and approvals as a source of implementation delays, governance is increasingly becoming the enabler or blocker of successful AI deployment.



Gareth Oldale

Partner and Head of Digital, Data and Commercial, TLT

Case study | Moving beyond pilots: GLD & Project GRAFO

The Government Legal Department (GLD) is exploring how AI can enhance the delivery of legal services, particularly legislative drafting. Project GRAFO is a prototype AI system designed to draft statutory instruments using a specially trained open source LLM.

Developed as part of GLD's broader AI strategy, Project GRAFO is exploring how large language models can support legislative drafting, while maintaining accuracy, accountability and quality standards.



The challenge

The GLD team identified numerous opportunities for AI, but deployment has been at a measured pace due to teams being conscious of the potential risks: the cost of errors, data loss and low-quality outputs.

Legislative drafting is a strong example. It demands precision and strict adherence to legal conventions, which led many to assume the task would be too complex for AI to support safely.



The approach

Project GRAFO is a structured, staged approach to development and deployment. Working with an external partner, the team began with a simple test case to assess feasibility, before moving on to a more complex pre-existing statutory instrument (SI). Each stage was used to refine the model and assess the quality of outputs as proof of concept.

Crucially, GLD combined technical deployment with strong governance. Developed in a secure environment, the system was subject to legal assurance processes, ensuring outputs underwent appropriate human and legal review before use.



The impact

This approach allows GLD to test AI in a controlled way, focusing on defined use cases, existing safeguards and rigorous oversight.



It's important that organisations maintain the quality of their legal services and leverage the existing processes they already have to support that."

David Lewis, Deputy Director – AI Programme,
Government Legal Department

2. Governance needs to move at the speed of delivery

Governance is one of the clearest pressure points shown by our research.

Over half of public sector organisations (56%) say governance and approvals are a major cause of delays when implementing new policies or service improvements. Meanwhile, 46% see complex governance as a barrier to improving efficiency and reducing duplication, and 34% cite a lack of governance and assurance frameworks as a key obstacle to adoption.

Many organisations are under pressure to deliver change faster, but there is also recognition that new technologies require clear oversight, skilled teams and strong data foundations. Without these, risks increase and confidence falls.

That lack of confidence is evident in our findings. Just 9% are very confident they can meet legal and regulatory expectations, while 23% are not confident at all and 22% are unsure. As a result, almost half lack confidence or clarity about whether they can meet the standards expected of them. The challenge is finding a governance model that enables progress while providing the assurance needed for responsible deployment.



TLT Takeaway

The answer is not necessarily more governance. It is better-designed governance. The good news is that this is often not a case of starting from scratch. Existing processes for legal, regulatory, communications and other workflows are likely to already include built-in assurance steps. Determining whether these processes can be adapted for AI is an important starting point.

Responsible AI use in the public sector needs governance models that operate as a system of triage and escalation, rather than a single blunt approval gate. Lower-risk use cases should move more quickly through proportionate review. Higher-risk use cases, such as those affecting prioritisation, eligibility, compliance, case handling, enforcement or citizen outcomes, will need stronger legal, data, risk and assurance scrutiny at the outset.

This is where legal teams can play a broader strategic role. General counsel and legal leaders should help define the organisation's decision thresholds before deployment. They have a vital perspective on what kinds of use case trigger closer scrutiny, what evidence is needed before a tool goes live, where accountability sits and when human oversight must be essential rather than nominal.

The big governance question is not “have we approved this AI tool?” It is “have we built a defensible route to deployment?” That is a different and more demanding standard, one that requires governance to be embedded into delivery instead of bolted on around it. Rather than a brake on innovation, well-designed governance should accelerate adoption by creating clarity around ownership, accountability and decision-making.



Nicola Mead-Batten
Partner and Head of Regulatory

Research findings

How confident are you in meeting your legal requirements and regulatory governance expectations?



are either not confident or unsure

What are the main causes of delays when implementing new policies or service improvements?



say governance or approvals are a major cause of delays

What non-technology organisational barriers most affect your ability to improve efficiency and reduce duplication?




identify complex governance as a barrier



Leadership actions

Governance models should reflect how materially the particular use case could affect decisions, services or outcomes.

1. Clarify who owns delivery, who owns risk and who can approve progression.
2. Engage legal and risk teams early.
3. Do not leave defensibility questions until the end of the process.
4. Match controls to impact.
5. Establish regular review periods to assess AI model performance, risks and suitability.



Client view



The challenge is that it's not just the technology that will evolve. Policies will need to be agile and iterative too so that they can adapt over time."

Marie Heracleous
Chief Digital and Innovation Officer,
Metropolitan Police

3. Procuring AI is becoming the first line of AI control

Buying AI well is now a strategic and legal capability.

Procurement of AI tools to improve efficiency and cost effectiveness has become one of the most important control points in public sector AI. Most organisations expect to rely heavily on third-party suppliers (57%) as they build AI capability. Currently, though, teams are facing challenges around limited internal expertise (53%), slow procurement cycles (48%), lack of AI-specific frameworks (47%), budget limitations (44%) and difficulty evaluating suppliers (40%).

The pressure on organisations to implement AI fast has turned procurement into a potential point of vulnerability as well as a sourcing challenge. If organisations lack the expertise or frameworks to interrogate supplier claims properly, they could get locked into poorly understood tools, insufficient transparency, weak audit rights or unclear accountability. This could not only slow down internal innovation, but damage service provision and cause serious issues if something goes wrong later.



TLT Takeaway

The quality of an AI deployment is significantly determined before a solution goes live. The decisions made at the point of supplier selection shape transparency, accountability, auditability and resilience throughout the entire deployment. Procurement is an important first forum where legal, commercial, governance and delivery questions meet.

These considerations in the procurement process are necessary because AI is not simply a software purchase. Traditionally, an organisation selects a solution, configures it, trains users, and manages it through periodic updates. Once deployed, the software behaves consistently according to predefined rules. AI systems generate responses based on probabilities and outputs which can vary depending on context, prompts, data sources, model updates and user behaviour. As a result, organisations need ongoing monitoring and validation rather than assuming performance will remain consistent over time.

Public sector organisations need to understand:

- What the system is doing.
- How it has been trained or configured.
- What evidence exists around performance.
- Where model limitations sit.
- How outputs can be interrogated.
- What recourse the organisation has if performance degrades or risk crystallises.

Procurement is one of the earliest AI governance controls. Contracts and supplier due diligence should assess explainability, resilience, data handling and accountability. Done well, procurement reduces risk, strengthens assurance and gives leaders confidence that deployment decisions can withstand scrutiny.




Vanessa Cooper

Partner
Digital, Data and Commercial

Research findings

What challenges do you face procuring AI for your organisation?



Leadership actions

Good AI procurement should test:

1. What exactly is the tool doing, and what is it not doing?
2. What data does it rely on, and how is that data handled?
3. What evidence supports performance, assurance and testing?
4. What level of explainability or auditability is possible?
5. What happens when performance degrades, or a risk issue emerges?
6. What rights exist to monitor, challenge, remediate or exit?



Client view



There are two instincts with AI. One is to push ahead because it's the future. The other is to hold back until all risks have been managed. The challenge is finding the middle ground between innovation and control.

Peter Shirley


Commercial Specialist, AI, Department for Work and Pensions

4. Trust depends on defensibility, not messaging

Trust is often discussed as a communications issue, but our research points to something more practical.

The leading concerns impacting trust are data governance (57%), explainability of AI decisions (51%), bias and unfair outcomes (44%), legal accountability and liability (44%) and safety and security risks (41%). These risks are set to become ever more pertinent as AI moves closer to everyday decision making. Organisations will need to show that using the outputs was reasonable, proportionate and appropriately governed.

What concerns do you have about compliance and accountability in AI?



TLT Takeaway

For legal teams, the trust question is fundamentally an evidential one. Organisations must prove that a use case complies with relevant legal and regulatory requirements from the start.

They must also maintain a credible record explaining why that use case was chosen, how potential bias was assessed and mitigated, and what controls or limits were set. Finally, they need to demonstrate how human oversight worked and how the AI's decisions can be challenged or reviewed after the fact.

This becomes more acute where AI begins to influence day-to-day decisions. Only 36% of organisations currently use AI to support operational decisions.

Once AI starts shaping prioritisation, risk scoring, resource allocation, compliance analysis or case handling, the defensibility burden rises materially.

Trust is therefore not merely about messaging, it is a matter of defensibility. Public confidence matters, but so does internal confidence too. If senior leaders, legal teams and delivery teams cannot see a clear evidence chain around a deployment, the organisation will struggle to scale even if the use case has clear value.




Jack Doukov-Eustice

Legal Director

Digital, Data and Commercial

Research findings

Which risks are the highest priority concerns for your organisation?



Leadership actions

For each significant use case, document:

1. Purpose and intended scope.
2. Data sources and control.
3. Bias assessment and mitigation measures.
4. Testing and assurance activity.
5. Explainability and documentation.
6. Human oversight points.
7. Challenge and review routes.
8. Audit trail and decision record.



Client view



Trust needs to be addressed through openness and transparency. We need to be able to show the benefits and rigour of our processes.”

Marie Heracleous
Chief Digital and Innovation Officer,
Metropolitan Police

5. Resilience is part of the licence to operate

AI failure should be expected, planned for and contained.

However, the research shows that operational resilience is one of the least mature areas of AI adoption in the public sector. Only 7% of respondents think their organisation is well prepared to respond to AI-related failures. 14% say they are not prepared, 23% are slightly prepared and 35% are only partially prepared. That means 72% are not yet mostly or well prepared for AI-related failure.

At the same time, 32% identify cybersecurity and safety as a priority AI risk, 17% identify operational failures, and 16% identify reputational risk. AI failure rarely stays technical for long in a public sector context. A poor output, model error, control failure or supplier issue can quickly become operational, reputational, political or legal.



TLT Takeaway

Failure readiness should be a precondition for scale, not something to develop once deployments are already live. Public sector organisations already understand resilience in other contexts: cyber incidents, outages, data breaches, supplier failure and service interruption. AI needs to be approached in the same disciplined way.

Leaders should assume that at some point an AI-enabled process will produce a problematic output, fail to perform as expected or surface governance issues that require escalation. What happens next needs to be clearly defined now.

From a legal perspective, this is critical. Failure readiness is part of the organisation's licence to operate AI in public service environments. If an organisation cannot identify what failure looks like, who owns escalation, when deployment should be paused and how review decisions will be evidenced - it is not yet ready for meaningful scale. Resilience is a core test of whether AI is ready for live deployment.

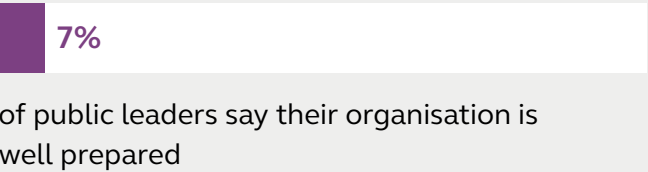


Anthony Vittadini

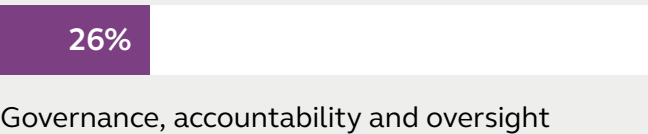


Legal Director
Digital, Data and Commercial

Research findings

How well prepared are you to respond to AI-related failures?



Which AI-related risks are the highest priority concerns for your organisation?



Leadership actions

Before live deployment, define:

1. Failure thresholds and trigger points.
2. Accountable owners and escalation routes.
3. Rollback or suspension steps.
4. Legal, risk and communications pathways.
5. Audit, review and lessons learned processes.



Client view



We don't yet truly understand how resilient existing regulatory and governance frameworks are for incorporating AI. The core issue is getting more clarity around who is responsible when things go wrong."

Melanie Venables
Director of Policy and Communication,
Professional Standards Authority

6. Workforce capability will determine who delivers real value

The public sector is preparing for AI-driven change in the workforce, but the data suggests the real gap is not awareness it is readiness.

Public sector leaders are clear on the scale of disruption ahead. 66% expect AI to drive significant or transformational change to workforce roles and skills, while only 8% expect minimal or no change.

However, the response does not yet match that shift. Most organisations are focusing on AI literacy (62%), practical AI skills (51%) and ethics, safety and governance (44%), with far less emphasis on specialist data capability (35%), process redesign (33%) or leadership for an AI-enabled workplace (22%).

This imbalance is important. The workforce challenge is not simply about equipping people to use AI tools. It is about reshaping how work is delivered, how decisions are made and where accountability sits. At the same time, 45% already identify skills


and capacity as a barrier to adoption, yet investment remains weighted towards foundational training rather than operational capability.

There is also a disconnect in risk perception. Only 16% identify workforce capability as a top AI-related risk, despite the scale of change expected. This suggests organisations recognise the impact of AI on roles, but have not yet translated that into a sufficiently robust workforce strategy. The implication is clear. Organisations are investing in skills, but not yet at the level of operating model, leadership and delivery redesign required to scale AI effectively.


That creates a structural constraint. AI will not deliver sustained value if it is layered onto existing ways of working without reshaping the capability, accountability and control environment around it.

Research findings


How much do you think AI will change workforce roles and skills?



What is the main barrier to AI adoption?



What training or reskilling do you plan?





TLT Takeaway

The greatest opportunity of AI is to reshape how work is delivered to enable better outcomes, not just deliver existing work faster. Whilst many teams are looking at individual AI literacy, far fewer at this stage are focusing on process redesign.

Our approach to AI Maturity identifies five stages of AI development: AI Readiness, Foundation, Adoption, Optimisation, Transformation, and sets out the use cases and organisational maturity that characterise each stage. Most public sector organisations currently sit at the Foundation or early Adoption stage - testing and experimenting with AI tools, and beginning to drive broader uptake through training and communications. Progressing beyond Adoption requires more than individual AI literacy. It demands the process redesign, accountability structures and leadership capability that our research shows are currently underdeveloped.

Workforce strategy needs to be calibrated to where an organisation sits on that maturity curve, and designed to close the gap to the next stage - not just equip people to use the tools they already have.

From an employment perspective, three priorities will shape what comes next:

1. Map current maturity and close the gap

Identify where the organisation sits and target activities which aim to close the gap needed to reach the next stage.

2. Redesign roles and accountability for the stage ahead

Clarify decision-making authority, redefine oversight responsibilities and restructure roles to reflect how AI-enabled work is actually delivered.

3. Build the capability for legal defensibility and scale

Equip staff not just to use AI tools, but to challenge outputs, justify decisions and evidence outcomes under scrutiny.



Amy McConnell
Partner, FutureLaw

Client view



We are creating Copilot communities across our workforce to help colleagues share good practice and identify emerging risks. Rather than banning public AI tools outright, we focus on responsible use, ensuring colleagues understand the guidance and never input sensitive or confidential information. We see this as a more effective way of combating shadow AI.”

David Lewis

Deputy Director AI Programme, Government Legal Department

7. The next wave: autonomy and exposure


The next wave has the potential to intensify existing weaknesses.

Looking towards future use cases for AI, 39% of public sector organisations identify predictive analytics and forecasting as the trend most likely to shape their operations. 27% point to cloud-based shared data platforms and interoperability. A further 24% identify automated casework, 24% identify AI cybersecurity platforms, and 21% point to agentic AI for autonomous decision-making.

Yet readiness remains low. Only 22% consider their organisation prepared or very prepared for the shift from traditional automation to agentic AI. 17% say they are not prepared, 26% are aware but unprepared, and only 4% say they are very prepared.

Without the right governance and guardrails in place, more autonomous systems are likely to amplify existing pressure points and in some cases, introduce new ones. If organisations already struggle with governance speed, explainability, supplier assurance, accountability mapping and

incident readiness, those issues will become more acute as AI systems become more embedded and able to act with less human oversight. Traditional controls may not continue to be effective as the technology evolves, and governance approaches need to take this evolution into account.



say their organisation is very prepared for the shift to agentic AI



The prevalence of AI over the coming years means that regulators are going to have to work together in a more cohesive way to drive forward change.”

Melanie Venables,
Director of Policy and Communication,
Professional Standards Authority



TLT Takeaway

Building solid foundations for successful AI adoption across the public sector requires sound legal and institutional design as well as the right tools and implementation partners.

Boards, executive teams and legal teams should prepare for a world in which accountability is more complex, system behaviour is less linear and conventional approval approaches don't provide enough assurance on their own. That means using today's deployments to build governance disciplines now, while the stakes are still manageable.



Emma Erskine-Fox
Partner
Digital, Data and Commercial

Conclusion

The era of AI experimentation is giving way to the era of implementation. For public sector leaders, the priority now is not finding more use cases, but creating the conditions for AI to deliver real and lasting value. That means strengthening governance, procurement, leadership and organisational readiness. These five actions highlight where organisations should focus next.



1. Prioritise fewer use cases with clearer value:

Focus on a small number of use cases linked to tangible strategic outcomes such as productivity, service improvement, decision quality, resilience or risk reduction. Avoid expanding pilot activity where there is no realistic route to operationalisation.



2. Build governance that enables trust:

Use a practical stage-gate process that classifies use cases by risk and impact, assigns ownership, and sets evidence requirements before deployment. Embed accountability, transparency and human oversight from the outset to build trust among leaders, employees and service users.



5. Invest in leadership and workforce readiness:

AI literacy across the organisation is important, but shouldn't come at the expense of preparing leaders to manage an AI-enabled workforce.

Leaders should be equipped to manage change, set direction and make informed decisions about AI use. Embedding AI understanding, accountability and cultural change at senior levels can help close the gap between expected workforce transformation and leadership readiness.



3. Treat procurement as part of assurance:

Build supplier evaluation and contract design around model governance, explainability, auditability, resilience, retraining, data handling and exit rights as well as functionality and price.



4. Define what failure looks like before going live:

Set out failure thresholds, pause or rollback triggers, escalation routes, communications plans and review processes before scaling deployment.

How TLT can support

We have experience advising central government departments, regulators, blue-light organisations and local authorities on AI governance, procurement, data management, automated decision-making, risk and public trust.

Our teams work directly with senior AI and legal leaders across government departments, helping to translate policy into action. Sharing our knowledge and forward views with our public sector clients can help them to deliver essential services in the most effective way with the greatest social impact.

We know that the scale, complexity and scrutiny attached to public service delivery is a unique challenge, and one that requires collaboration and common purpose to navigate. Our team has a depth of understanding on how the technology works as well as the various stages of an AI project lifecycle and the risks at each stage.

We regularly advise clients on their AI adoption journey by helping them evaluate tech tools, run pilots, build internal business cases, provide training, plan communications, and develop use cases.

visit [**tlt.com**](https://www.tlt.com)



Gareth Oldale
Partner

London
+44 (0)7584 706 566
[**gareth.oldale@tlt.com**](mailto:gareth.oldale@tlt.com)



Emma Erskine-Fox
Partner

Bristol
+44 (0)7811 805 133
[**emma.erskine-fox@tlt.com**](mailto:emma.erskine-fox@tlt.com)



Amy McConnell
Partner

London
+44 (0)7483 399 027
[**amy.mccconnell@tlt.com**](mailto:amy.mccconnell@tlt.com)



Vanessa Cooper
Partner

London
+44 (0)7976 517 827
[**vanessa.cooper@tlt.com**](mailto:vanessa.cooper@tlt.com)



Anthony Vittadini
Legal Director

London
+44 (0)7795 120 884
[**anthony.vittadini@tlt.com**](mailto:anthony.vittadini@tlt.com)



Jack Doukov-Eustice
Legal Director

London
+44 (0)7788 364 498
[**jack.doukov-eustice@tlt.com**](mailto:jack.doukov-eustice@tlt.com)

AI glossary of terms

Agentic AI: AI systems composed of agents that can behave and interact autonomously to take actions, make decisions, and complete multi-step tasks to achieve their objectives.

Why it matters: Introduces a step-change in autonomy, increasing both productivity potential and governance risk.

AI (Artificial Intelligence): Systems that perform tasks typically requiring human intelligence, such as decision-making, prediction, language processing, or pattern recognition.

Why it matters: AI is now embedded in operational decisions, not just analytics or automation.

AI Lifecycle: The end-to-end stages of an AI system: design, data preparation, development, testing, deployment, monitoring, and retirement.

Why it matters: Governance and legal accountability must span the full lifecycle, not just deployment.

Algorithmic Bias: Systematic errors in AI outputs that lead to unfair or skewed outcomes for certain groups.

Why it matters: A key legal and reputational risk area, particularly in public service delivery.

Automated Decision-Making (ADM): When decisions are made by systems without meaningful human intervention.

Why it matters: Highly sensitive in public sector contexts due to legal, fairness, and accountability implications.

Data Governance: The policies and controls that define how data is collected, stored, shared, and used.

Why it matters: Poor data governance is one of the most common blockers to safe AI deployment.

Explainability: The ability to understand and articulate how an AI system reached a decision or output.

Why it matters: Essential for legal defensibility, public trust, and regulatory compliance.

Generative AI (GenAI): AI systems that generate new content (text, images, code, summaries) based on instructions given to it by a user.

Why it matters: Widely used in public sector productivity tools, but increases risks around accuracy, bias, and data leakage.

Human-in-the-loop (HITL): A model where humans remain actively involved in reviewing, approving, or overriding AI outputs.

Why it matters: A critical safeguard for high-risk decisions affecting citizens.

Large Language Model (LLM): A type of generative AI trained on vast text datasets to understand and generate human-like language.

Why it matters: Powers tools like chat-based assistants used in citizen services, drafting, and analysis.

Machine Learning (ML): A subset of AI where systems learn patterns from data to make predictions or decisions without being explicitly programmed.

Why it matters: Often sits behind operational tools such as fraud detection, forecasting, and risk scoring.

Model Drift: When an AI system's performance degrades over time because real-world data changes from the data it was trained on.

Why it matters: A key operational risk that can quietly undermine accuracy and reliability.

Model Governance: The frameworks, controls, and oversight used to manage AI across its lifecycle (design, deployment, monitoring, retirement).

Why it matters: Core requirement for defensible and compliant AI use in regulated environments.

Prompt (in GenAI systems): The instruction or input given to a generative AI system to produce an output.

Why it matters: Small changes in prompts can materially change outputs, raising consistency and control challenges.



tlt.com

Belfast | Birmingham | Bristol | Edinburgh | Glasgow | London | Manchester | Piraeus

TLT LLP and TLT NI LLP (a separate practice in Northern Ireland) operate under the TLT brand and are together known as 'TLT'. Any reference in this communication or its attachments to 'TLT' is to be construed as a reference to the TLT entity based in the jurisdiction where the advice is being given. TLT LLP is a limited liability partnership registered in England & Wales number OC308658 whose registered office is at One Redcliff Street, Bristol, BS1 6TP. TLT LLP is authorised and regulated by the Solicitors Regulation Authority under ID 406297.

In Scotland TLT LLP is a multinational practice regulated by the Law Society of Scotland.

TLT (NI) LLP is a limited liability partnership registered in Northern Ireland under ref NC000856 whose registered office is at River House, 48-60 High Street, Belfast, BT1 2BE

TLT (NI) LLP is regulated by the Law Society of Northern Ireland under ref 9330.

TLT LLP is authorised and regulated by the Financial Conduct Authority under reference number FRN 780419. TLT (NI) LLP is authorised and regulated by the Financial Conduct Authority under reference number 807372. Details of our FCA permissions can be found on the Financial Services Register at <https://register.fca.org.uk>

Methodology

This report is based on research conducted in May 2026, combining a survey of more than 100 public sector leaders and one-to-one interviews across central government, local government, regulators and blue-light organisations. The findings provide insight into current views and experiences of AI adoption across the sector. This report is for general information only and should not be relied upon as legal or professional advice.