DELIVERED BY: BYTES | Microsoft | GOV.NEWS

Shadow AI in the UK Public Sector:

From Unapproved Use to Safe, Strategic AI

Table of Contents

Connect with us

Executive Summary

The public sector is losing control of its sensitive data. 71% of employees use AI outside of organisationally approved systems, with 51% doing so weekly.[1] Well-intentioned staff may enter sensitive citizen or organisational information into consumer AI tools without realising the consequences.

Every unauthorised AI interaction creates a potential data breach, GDPR violation or reputational crisis. The extent of the problem is often invisible to IT teams - and senior leadership.

Bytes believes the only sustainable response to Shadow AI is strategic enablement, not suppression. Blanket bans will only drive the problem underground.

As a leading UK Microsoft partner with 40+ years public sector experience, Bytes has developed the Shadow AI 2026 framework to guide organisations from unmanaged risk to safe, enterprise-grade AI adoption. This five-pillar approach helps organisations strengthen leadership ownership, strengthen data foundations, deploy approved tools including Microsoft 365 Copilot, implement clear governance, and build cyber resilience.

The goal is to work with human behaviour rather than against it. Bytes is moving the conversation from technology towards leadership, operating models and capability gaps.

Why Shadow AI, Why Now

AI has become integral to daily life. In the government's 2026 'AI Skills for Life and Work' survey, 73% of the public used AI in day-to-day life in the past month, including passive tools like predictive text, virtual assistants, generative AI, and chatbots.

The proportion feeling confident using AI tools in their daily lives has increased compared to earlier studies, underlining how quickly AI use has expanded from a niche to a routine activity.

AI hype has resulted in rising expectations. The public anticipates rapid workplace integration, with 51% of workers predicting routine use within a year.[2] There's strong appetite to use AI to make tasks faster and easier.

At the same time, public sector policy is championing AI. The government-backed *AI Opportunities Action Plan* articulates the transformative potential of AI in the public sector.[3] The drive to adopt AI spans sectors. The *NHS 10-Year Health Plan* sets an ambition for developing the 'most AI-enabled health system in the world,' while the *Defence AI Playbook* suggests the potential for AI to secure the UK a strategic advantage.[4] [5] AI is seen as a way to tackle service backlogs, drive efficiency and enhance outcomes.

The message is clear - public sector workers are expected to embrace AI. In a sector under pressure to do more with less, people are looking to AI for the answers.

While policy actively encourages AI adoption, Bytes believes that the operating models and governance systems needed to support safe, strategic use have not kept pace.

The policy landscape signals ambition and expectation, but few organisations have the enterprise-grade infrastructure, clear policies, or workforce capabilities in place to deliver on those expectations safely.

Bytes' experience across the public sector reveals a consistent pattern. Shadow AI is not a technology failure, but an organisational readiness gap. Enterprise-grade AI like Microsoft CoPilot isn't always available to public sector workers. A Deloitte survey found that nine in ten organisations lack a policy on GenAI use and a governance structure.[6] AI adoption is being encouraged at scale, while governance, capability, and control mechanisms lag behind.

Shadow AI thrives in this gap between demand and sanctioned deployment. Naturally, staff are turning to familiar, easily available consumer tools in the absence of workplace solutions.

Leaders must close this readiness gap to fully harness the potential of AI while safely managing the risks.

The Reality of Shadow AI in the Public Sector

Shadow AI is already widespread. A Microsoft survey reveals that 71% of UK employees (including the public sector) have used unapproved consumer AI tools at work, with 51% doing so each week, far higher than many leaders anticipate.[7]

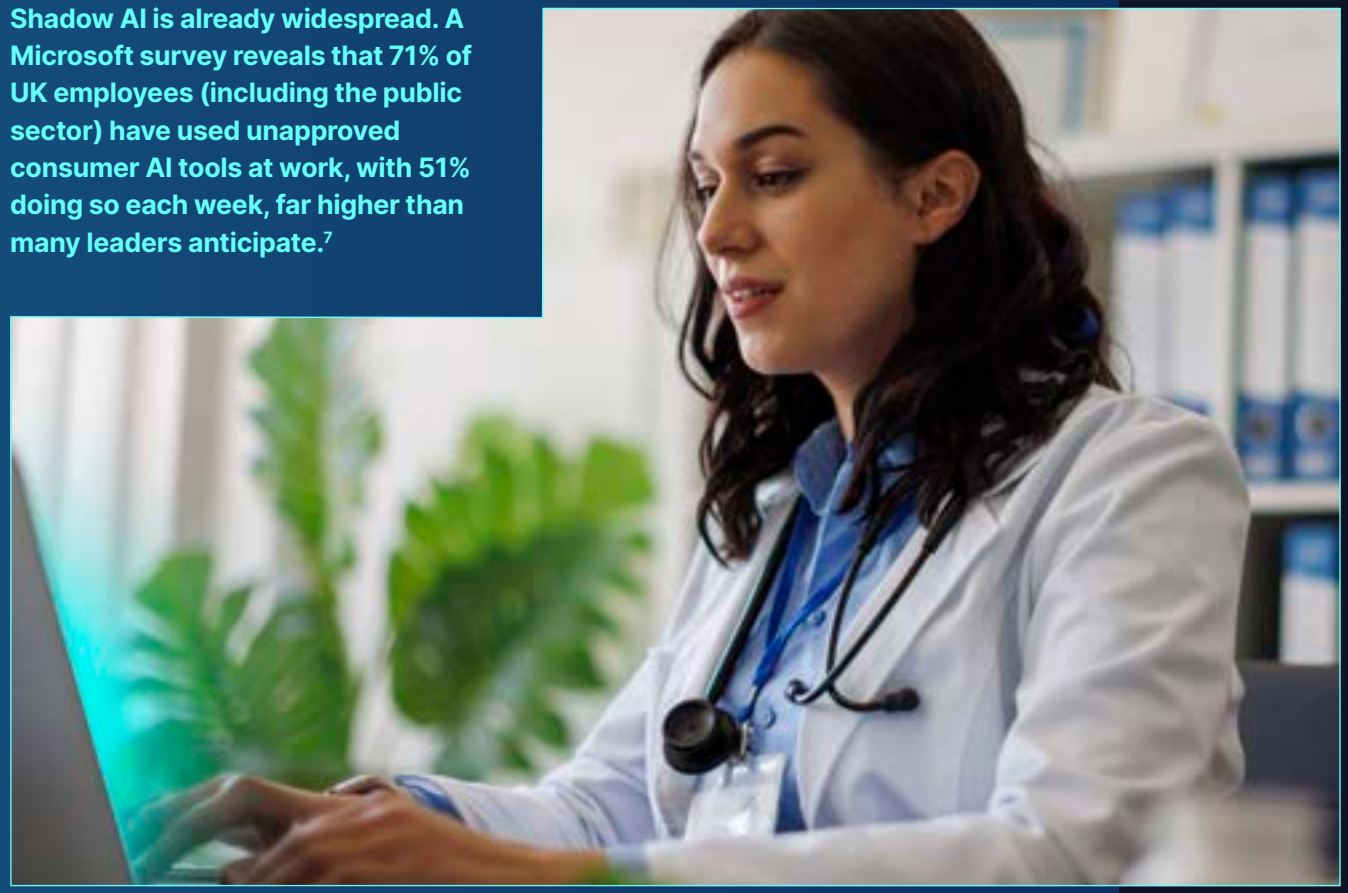

AI use isn't confined to junior staff. The report reveals use is consistent across seniority levels, with managers and executives tuning to personal AI tools for drafting reports and strategic communications.

The pattern spans sectors. Whether it's health and social care or education, staff are turning to familiar tools for drafting communications (49%), writing reports (40%) and analysing finances (22%) due to absent or unfamiliar enterprise options. AI use spans roles, and is most common among employees in IT, finance and marketing teams.

In healthcare, a 2024 Harvard study found that 16 per cent of GPs admitted to using ChatGPT to help treat patients, raising concerns about patient confidentiality and data protection.[8]

Some organisations have resorted to blanket bans to deal with shadow AI. For example, The Department for Work and Pensions explicitly banned ChatGPT in 2024, though later relaxed the policy.[9] Bytes believes that the scale of unapproved AI use shows it is systemic behaviour, not isolated misconduct.

Why Well-intentioned Staff Bypass Official Systems

Understanding why people resort to shadow AI helps leaders tackle the root problems. Several reasons drive the behaviour:

Lack of sanctioned systems
Without an approved organisational system, many turn to unapproved alternatives.

Slow approval processes
Requests for official systems can often take a long time to go through the relevant governance processes, pushing people elsewhere.

Low awareness of risk
The Microsoft survey revealed that many staff didn't fully appreciate the risks of using consumer tools, signalling a genuine lack of understanding.

Cultural silence
When leaders don't openly discuss AI use, staff have no framework for distinguishing safe from unsafe practices and determining acceptable use.

Approved tools don't meet needs
Where sanctioned tools exist, they may lack the functionality, familiarity, simplicity, or responsiveness of consumer alternatives.

Shadow AI is a symptom, not the root cause. Bytes' work with public sector organisations reveals that staff aren't looking for workarounds - they're looking for tools that work.

This is why Bytes partners with Microsoft to deliver solutions like Microsoft 365 Copilot: consumer-grade simplicity with enterprise-grade security.

Bytes believes the real issue is a lag in organisational readiness rather than individual rule-breaking. This is why blanket bans don't work. Understandable reasons drive the behaviour. Guidance without alternatives will be ignored because the use of unsanctioned tools is fulfilling an unmet need.

Enforcement consistently lags behind behaviour. By the time organisations detect and respond to shadow AI, new tools and workarounds have already emerged.

Organisational policy needs to work with the human element, not against it. Punishing well-intentioned behaviour risks driving shadow AI underground. Enforcing a workplace ban will see people turn to personal devices for work use. This creates a much more insidious problem, making it almost impossible for organisations to monitor and address.

Bytes is moving the conversation away from policy breaches toward leadership, organisational readiness and capability gaps.

Why Shadow AI Is a Public Sector Risk Multiplier

Shadow AI brings hidden risks. At the point where an employee enters information into a public AI tool, that data is beyond the organisation's control. The AI supplier can use that data in accordance with its terms of use, including to train AI models.

This was highlighted in 2025, when ChatGPT conversations were indexed by Google after users unknowingly granted permission through sharing features, demonstrating how easily staff can create data exposure risks without realising the implications.[10]

Shadow AI multiplies cyber risk. Unseen data flows hinder an organisation's ability to minimise the risk of data breaches.

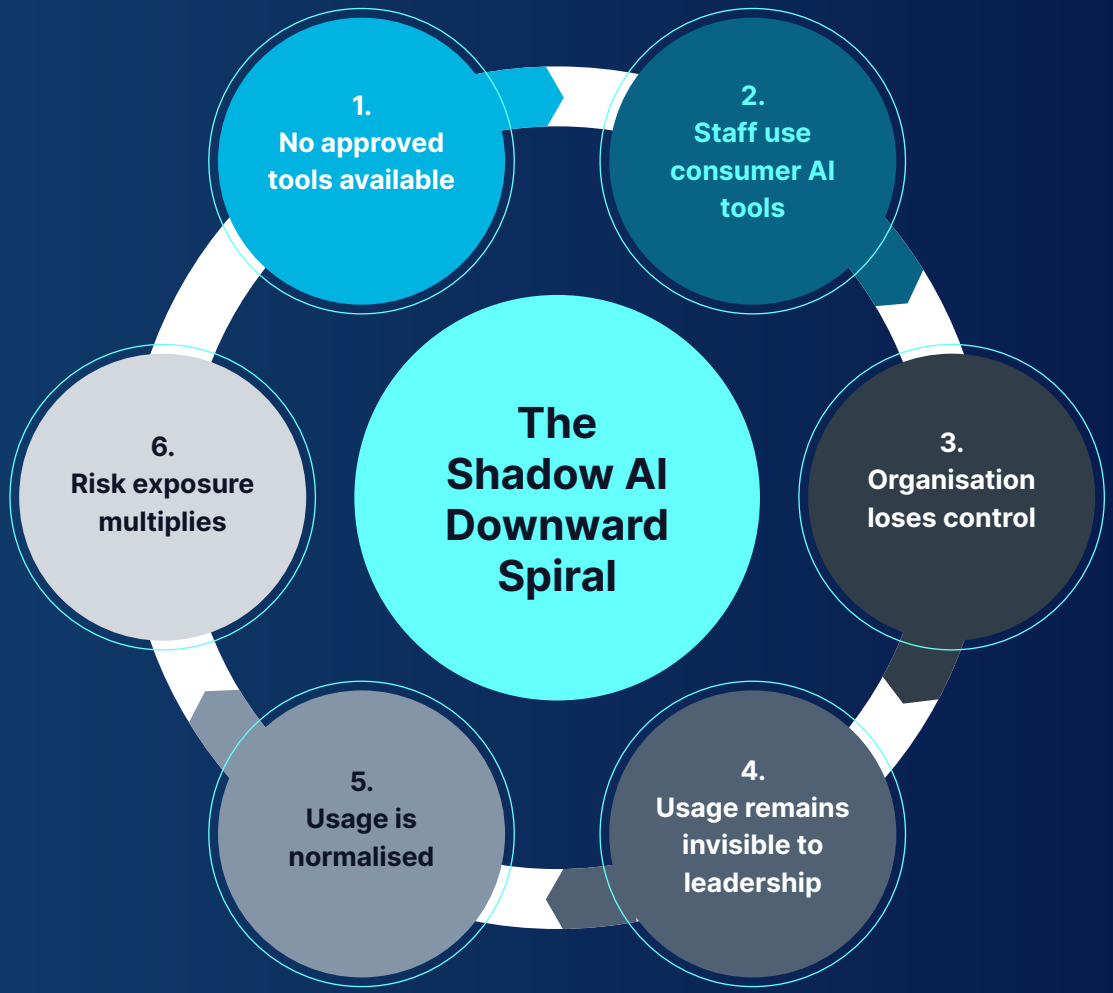

The stakes are high in the public sector, where staff are often handling sensitive citizen data, including safeguarding, health, and criminal records. Not fully aware of the risks, staff may enter this data into publicly available tools, breaching GDPR regulations and exposing the organisation to legal and reputational risks. In the defence sector, leaking information into public AI tools could pose a threat to national security.

Ultimately, shadow AI can impact people's lives. This was illustrated by a high-profile case in Australia. One social worker turned to a public AI tool to speed up writing reports on vulnerable children. Without realising, they had made sensitive, identifiable information available publicly, placing a vulnerable child at risk.[11]

Bytes' experiences shows that while the impact is serious, the workforce is largely unaware of the risks. Most inappropriate use is well-intentioned and inadvertent. For organisations, the level of exposure is largely invisible. Many organisations don't know where AI is being used today. Leaders must act now to take back control and minimise the risks.

Introducing the Shadow AI 2026 Framework

In response to the growing problem of shadow AI, Bytes Software Services has developed a five-pillar strategic operating framework to enable public sector organisations to move from unmanaged AI usage to safe, enterprise-grade adoption.

With over 40 years of experience, Bytes has grounded its Shadow AI 2026 framework in real-world experiences of supporting public sector organisations from HMRC to NHS England with responsible AI. Built specifically for public-sector complexity, the framework takes a nuanced approach, encompassing both control of shadow AI and enablement of approved use. Bytes believes that this dual response is the only way to drive sustainable behavioural change.

The framework features five interlocking pillars that must work together. No single control or policy solves shadow AI. The drivers are complex and require a multifaceted approach. Weakness in one pillar undermines the rest.

Fundamentally, shadow AI is a systemic problem. Bytes takes the view that it demands a whole system response, not a narrow focus on compliance.

> "Businesses must ensure the AI tools in use are built for the workplace, not just the living room. The message is clear: only enterprise-grade AI delivers the functionality that employees want, wrapped in the privacy and security every organisation demands."
>
> Darren Hardman, CEO, Microsoft UK & Ireland.[12]

The Five Pillars of Shadow AI 2026

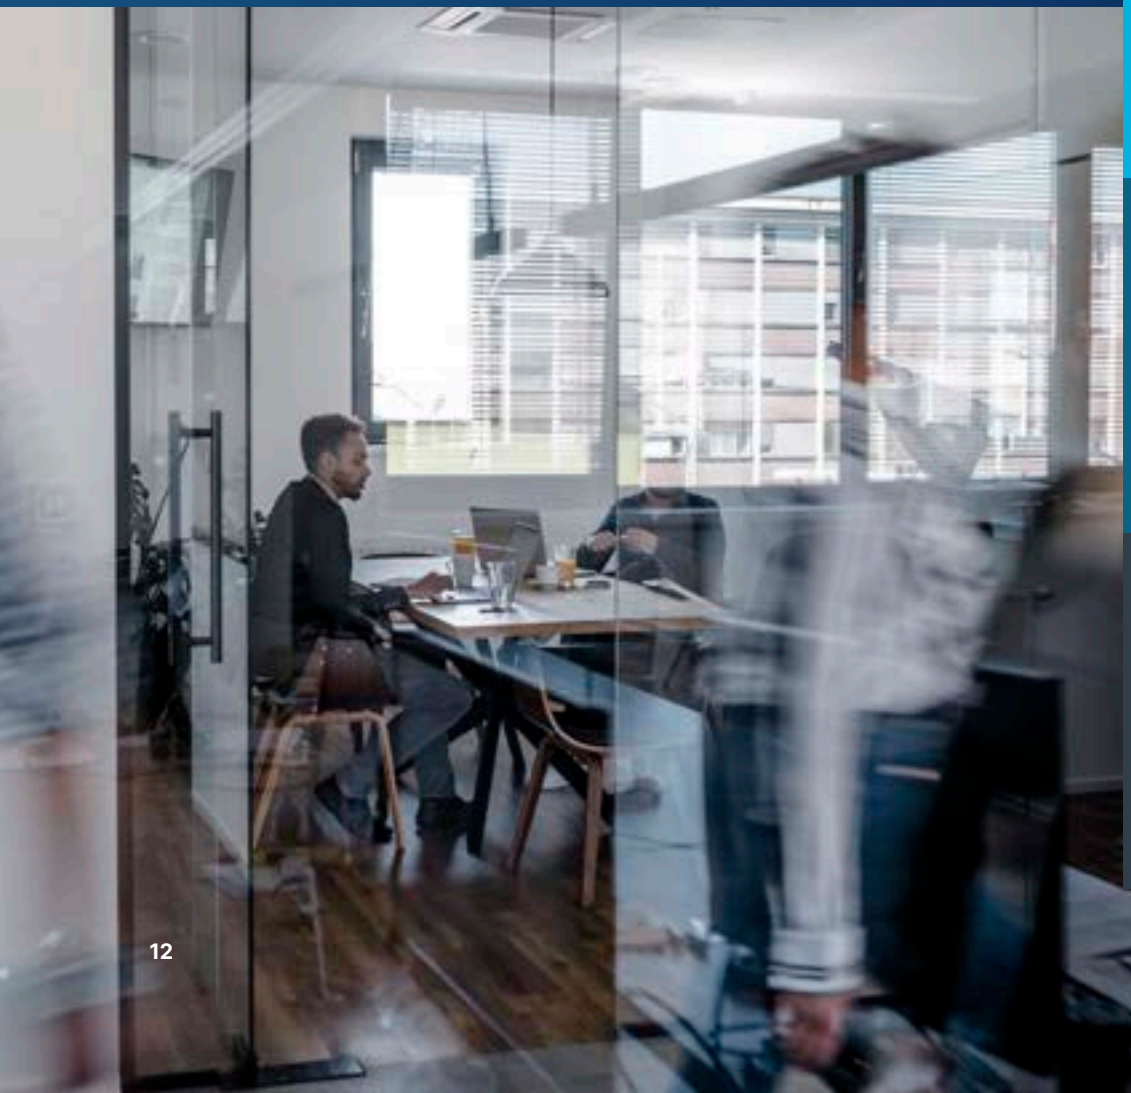

1

Leadership & Skills
Establishing executive ownership, AI fluency, and workforce capability to make informed decisions about AI use.

2

Data Foundations
Ensuring data is accessible, governed, and structured to support safe AI deployment.

3

AI Orchestration
Providing integrated, approved AI platforms that deliver the outcomes staff need without the unintended risks.

4

Governance & Control
Implementing clear policies, approval processes, and monitoring mechanisms that enable speed without sacrificing oversight.

5

Cyber, Risk & Resiliency
Mitigating the cyber and operational risks introduced by AI, including data leakage, model vulnerabilities, and insider threats.

Each pillar addresses a specific dimension of the shadow AI challenge. Together, they form a coherent strategy for organisations ready to take a nuanced response that doesn't just drive shadow AI underground.

The Five Strategic Pillars of Shadow AI 2026

Pillar 1: Leadership & Skills
Establishing executive ownership, AI fluency, and workforce capability to make informed decisions about AI use.

The problem

Shadow AI thrives in a leadership vacuum. Skills and knowledge gaps at the top of an organisation shape workplace culture.

In a Microsoft survey featuring respondents at all levels, 36% said they didn't know where to start with AI and a third weren't concerned about privacy and security risk.[13] This highlights skills and knowledge gaps across the board, with senior management attitudes shaping workplace culture.

When leaders fail to give clear direction about AI strategy and model acceptable use, decision-making defaults to individual staff making informal, and sometimes uninformed, judgements. Without leadership and skills, every other governance control breaks down.

Why this matters in the public sector

In a politically sensitive environment, accountability and public trust are paramount. Public sector organisations answer to citizens and regulators. Bytes' stance is that AI decisions require strategic oversight - they should not be confined to the IT team. Leadership gaps create operational, reputational, and political risk.

What good looks like...
The leadership and skills pillar is about ensuring the organisation has the clarity, confidence, and capability at every level to make informed decisions about AI.

Executive ownership
A named senior leader (e.g., Chief Data Officer, Chief Digital Officer) owns the AI strategy and is accountable for AI risk and opportunity.

AI fluency at the top
Senior leaders understand AI's capabilities, limitations, and risk profile, enabling them to make informed strategic decisions.

Workforce skills development
Staff at all levels receive training on AI literacy, prompt engineering, responsible use, and when to seek support.

Embedded strategy
The organisation's AI approach is documented, communicated, and reinforced through regular engagement.

Common failure patterns

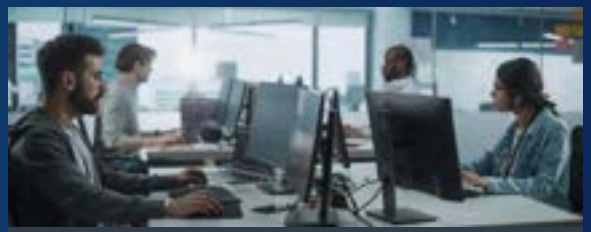

Delegating AI entirely to IT or cybersecurity teams without strategic oversight.

Assuming AI is 'someone else's problem' and not taking ownership.

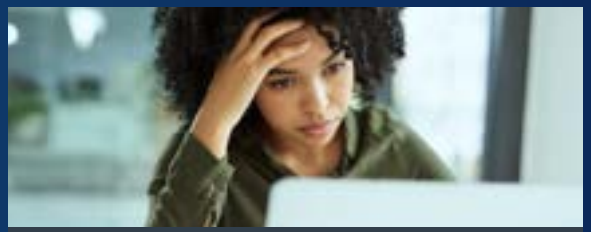

Failing to invest in AI skills at the top of the organisation, resulting in poorly informed decision-making.

Treating AI as a technical issue rather than a challenge requiring fundamental organisational change.

Strategic outcomes
When organisations master this pillar, leaders provide clarity and confidence about the use of AI. Staff know what is classed as acceptable and unacceptable use and understand how to use AI responsibly.

Pillar 2: Data Foundations
Ensuring data is accessible, governed, and structured to support safe AI deployment.

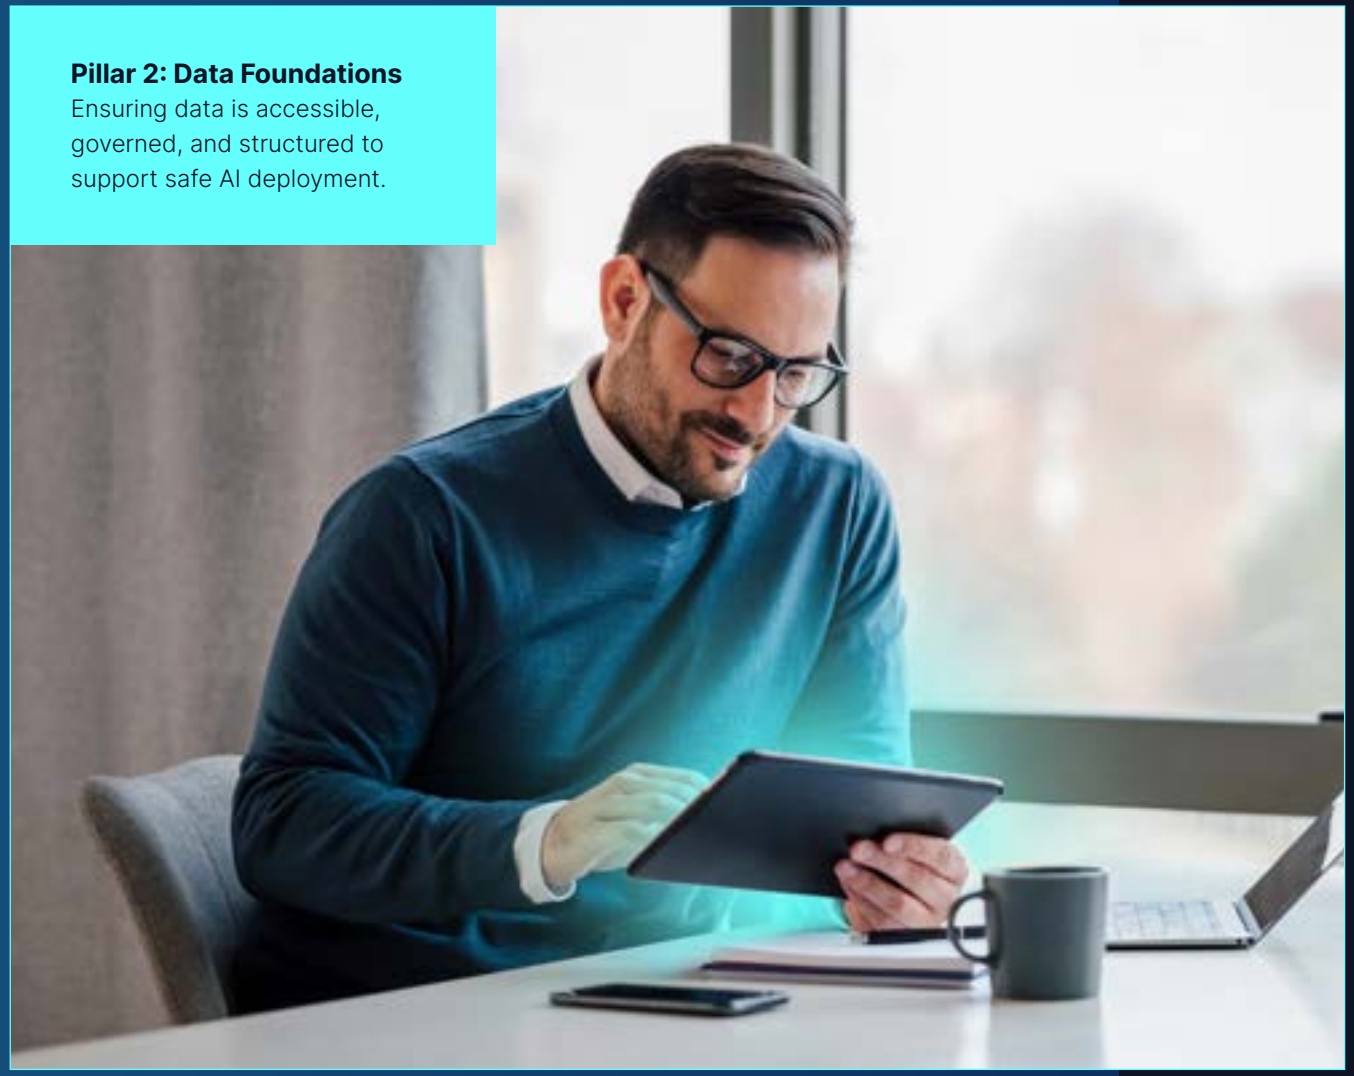

The problem

Shadow AI often emerges because sanctioned systems don't provide access to the data staff need. If critical information is locked in silos or difficult to extract, staff will resort to unsafe workarounds, including pasting data into external AI tools. Bytes' position is that good data is the precondition for safe AI.

AI maturity is capped by data maturity. Organisations cannot fully harness the benefits without the right data foundations in place.

Why this matters in the public sector

For public sector organisations, the stakes are high. Leaked information could put vulnerable citizens or even national security at risk. Reputational damage and political fallout are real threats. Data breaches can compromise safeguarding cases, expose health records, or undermine justice proceedings, carrying profound ethical and legal consequences.

What good looks like...
Data foundations is about ensuring the organisation's data is governed, accessible, and properly structured - creating the secure base on which safe AI can be built.

Robust data governance
Bytes recommends Microsoft Purview to automate sensitivity labelling and access controls.

Data is structured and accessible
Staff can find and use the information they need without resorting to unsanctioned tools.

Data quality is maintained
Information is accurate, up-to-date, and free from unnecessary duplication.

Permissions align with roles
Staff have access to the data they need - and only the data they need - based on their role and responsibilities.

Common failure patterns

Data is scattered across disconnected systems with no integration.

Sensitivity labels are missing or inconsistently applied.

Permissions are set too broadly, creating risk of over-exposure.

Stale, duplicated, or unstructured content can be surfaced inappropriately by AI tools.

Strategic outcomes
When organisations master this pillar, staff can access the information they need securely, reducing the temptation to use external tools. AI systems operate on well-governed, classified data, minimising the risk of accidental exposure. Data becomes an enabler of safe AI adoption, not a barrier.

Pillar 3: AI Orchestration
Providing integrated, approved AI platforms that deliver the outcomes staff need without the unintended risks.

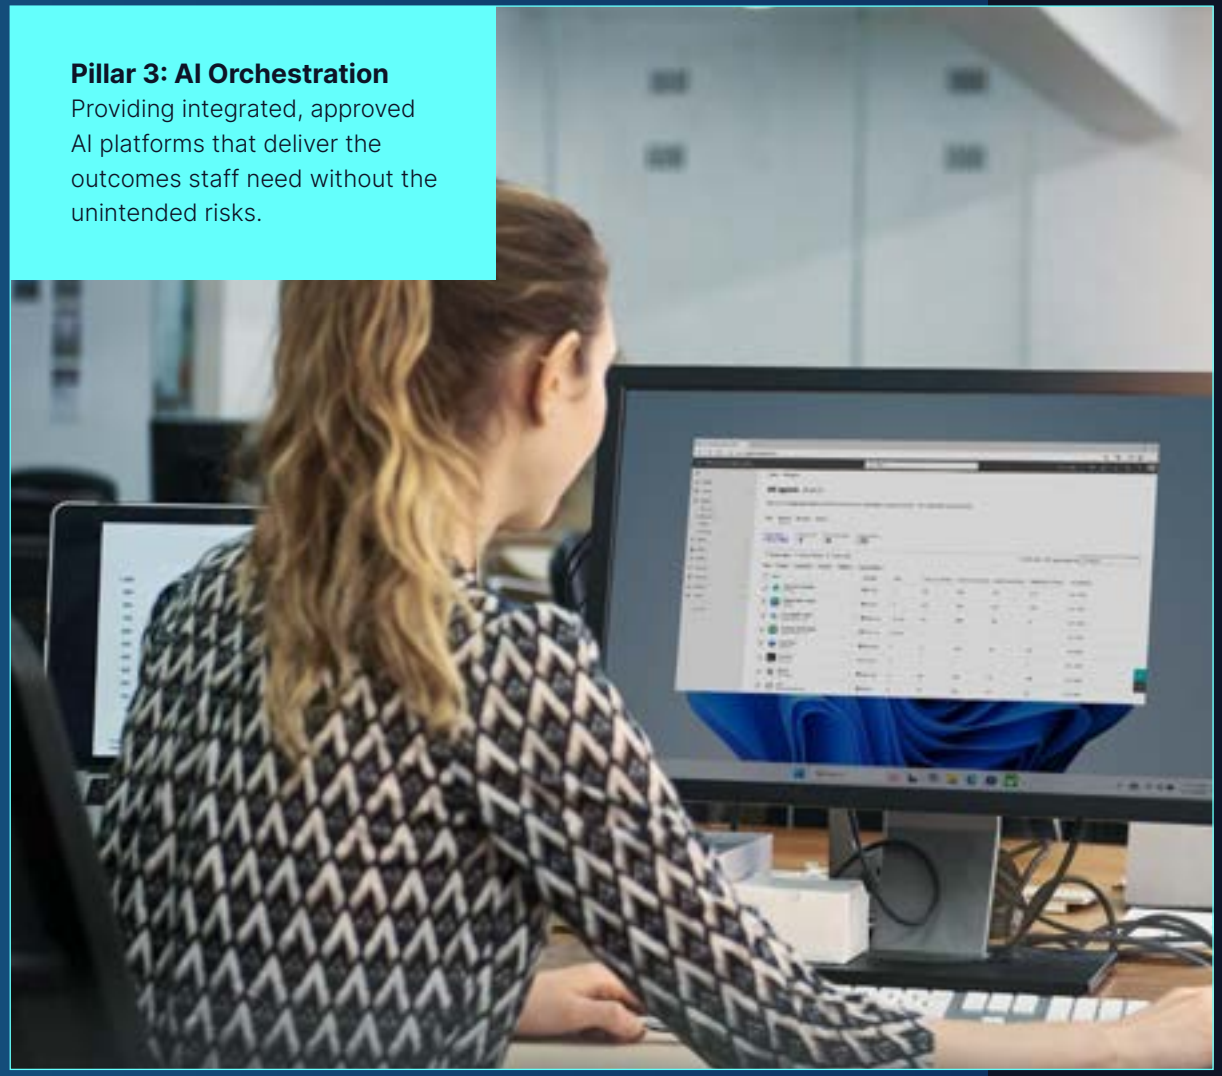

The problem

Staff are looking for outcomes, not specific tools. They need to summarise a 50-page consultation response, analyse the last quarter's service data or produce a first draft of a policy briefing. Bytes contends that when approved AI platforms are absent or don't immediately produce the desired outcome, individuals default to their preferred consumer tools. This ad-hoc approach means risk is invisible and organisations lose control of how AI is used.

As AI evolves, orchestration must extend beyond assistants to AI agents. Platforms such as **Microsoft Agent 365** enable organisations to manage agents across workflows, systems and data sources. Bytes believes that without central orchestration, this risks recreating shadow AI at greater speed and scale. Effective orchestration ensures that AI (including agents) is deployed within approved environments, aligned to defined use cases, and remains visible and controllable.

Why this matters in the public sector

Fragmented AI use equates to duplicated spend, compounding financial pressures. Worse, shadow AI means the organisation has no line of sight into how AI is being used and cannot stand up to public scrutiny and demonstrate value for money. Strategic orchestration is both a financial and a governance imperative.

What good looks like...
AI orchestration is about the organisation actively driving AI adoption - identifying priority use cases, providing integrated platforms, and enabling innovation while maintaining strategic control.

Sanctioned AI is embedded in daily workflows

Microsoft 365 Copilot enables orchestration. Integrated into Outlook, Word, Excel, and Teams, it makes AI assistance ubiquitous and controlled.

Data stays within organisational boundaries

AI processing happens within the Microsoft 365 tenant under enterprise data governance, not on external public platforms.

Use cases are clearly defined and communicated

Staff understand how to use AI for their everyday tasks - drafting emails, summarising meetings, and generating reports - without creating unnecessary risk.

AI delivers measurable value

Productivity gains are tracked, communicated, and celebrated, reinforcing adoption and demonstrating ROI.

Common failure patterns

Procuring AI tools but failing to integrate them into workflows.

Providing access without training based on everyday tasks, resulting in low adoption.

Choosing tools that don't align with existing platforms, creating friction.

Over-complicating deployment with excessive restrictions that undermine usability.

Strategic outcomes

When organisations master this pillar, staff have access to powerful, approved AI that meets their needs. Productivity increases visibly and measurably.

Shadow AI usage declines organically because the sanctioned alternative is better. The organisation maintains control, visibility, and compliance.

Pillar 4: Governance & Control
Implementing clear policies, approval processes, and monitoring mechanisms that enable speed without sacrificing oversight.

The problem

Bytes believes that ambiguity accelerates shadow AI. When staff don't know what's allowed and what's prohibited, they make their own judgements. Inconsistent rules, unclear policies, and absent monitoring create an environment where unsanctioned AI thrives. The Society for Technology, Innovation and Modernisation suggests that instead of saying 'no' and being ignored, public sector organisations should say 'yes, but safely' to AI.[14]

Why this matters in the public sector

Public sector organisations operate under strict regulatory frameworks and must stand up to public scrutiny. Yet staff will naturally circumvent governance that stifles innovation and slows productivity. Effective governance keeps pace with usage, securing compliance and earning public trust.

What good looks like...
Governance and control is about setting clear boundaries for AI use, monitoring how it's being used, and creating the policies and processes that keep AI adoption visible and accountable.

Clear, accessible AI usage policy
Staff know what AI tools are approved, what data can be used, and what actions are prohibited.

Approval processes are streamlined
Requests for new AI use cases or tools are reviewed quickly, balancing risk and opportunity.

Monitoring and audit trails
AI usage is logged, reviewed, and auditable. Organisations can answer 'who used AI, when, and for what purpose?'

Escalation pathways are defined
Staff know where to go if they encounter an AI-related issue, need guidance, or want to propose a new use case.

Common failure patterns

Policies that are vague, overly complex, or buried on an intranet site.

Approval processes so cumbersome that they incentivise circumvention.

No monitoring capability, leaving actual usage invisible.

Treating governance as a one-time policy document rather than an ongoing, adaptive practice.

Strategic outcomes
When organisations master this pillar, governance provides clarity and confidence. Staff operate within clear boundaries, knowing they're supported when they follow the rules. Monitoring provides visibility into usage patterns, enabling proactive risk management. Governance becomes an enabler of productivity, not a blocker.

Pillar 5: Cyber, Risk & Resiliency
Mitigating the cyber and operational risks introduced by AI, including data leakage, model vulnerabilities, and insider threats.

The problem

Shadow AI creates an invisible attack surface. When individuals use unapproved tools, they inadvertently introduce unseen data flows, making organisations vulnerable to cyber threats. Accidental leaks and insider risks multiply when AI usage is uncontrolled, increasing the risk of a disruptive and expensive cyberattack.

Why this matters in the public sector

Public sector organisations are likely to be the target of bad actors looking for intelligence, financial gain or disruption. A cyberattack can cause significant financial costs and threaten human lives or national security. As critical national infrastructure, cyber resilience is about maintaining the continuity of essential services. Bytes is clear that shadow AI is one of the greatest threats to public sector cybersecurity.

What good looks like...
Cyber, risk and resiliency is about protecting the organisation from the security threats that shadow AI creates, ensuring data stays where it should and vulnerabilities are identified before they're exploited.

Visibility into AI usage: Bytes implements
Microsoft Purview and Microsoft Defender for Cloud Apps make visible where unsanctioned AI is being used.

Data loss prevention controls
Automated policies detect and block attempts to send sensitive data to external AI platforms.

Insider risk monitoring
Anomalous behaviour, like mass file downloads and unusual access patterns, trigger alerts before data is exfiltrated.

Incident response readiness
The organisation has clear protocols for responding to AI-related data breaches, including containment, investigation, and notification.

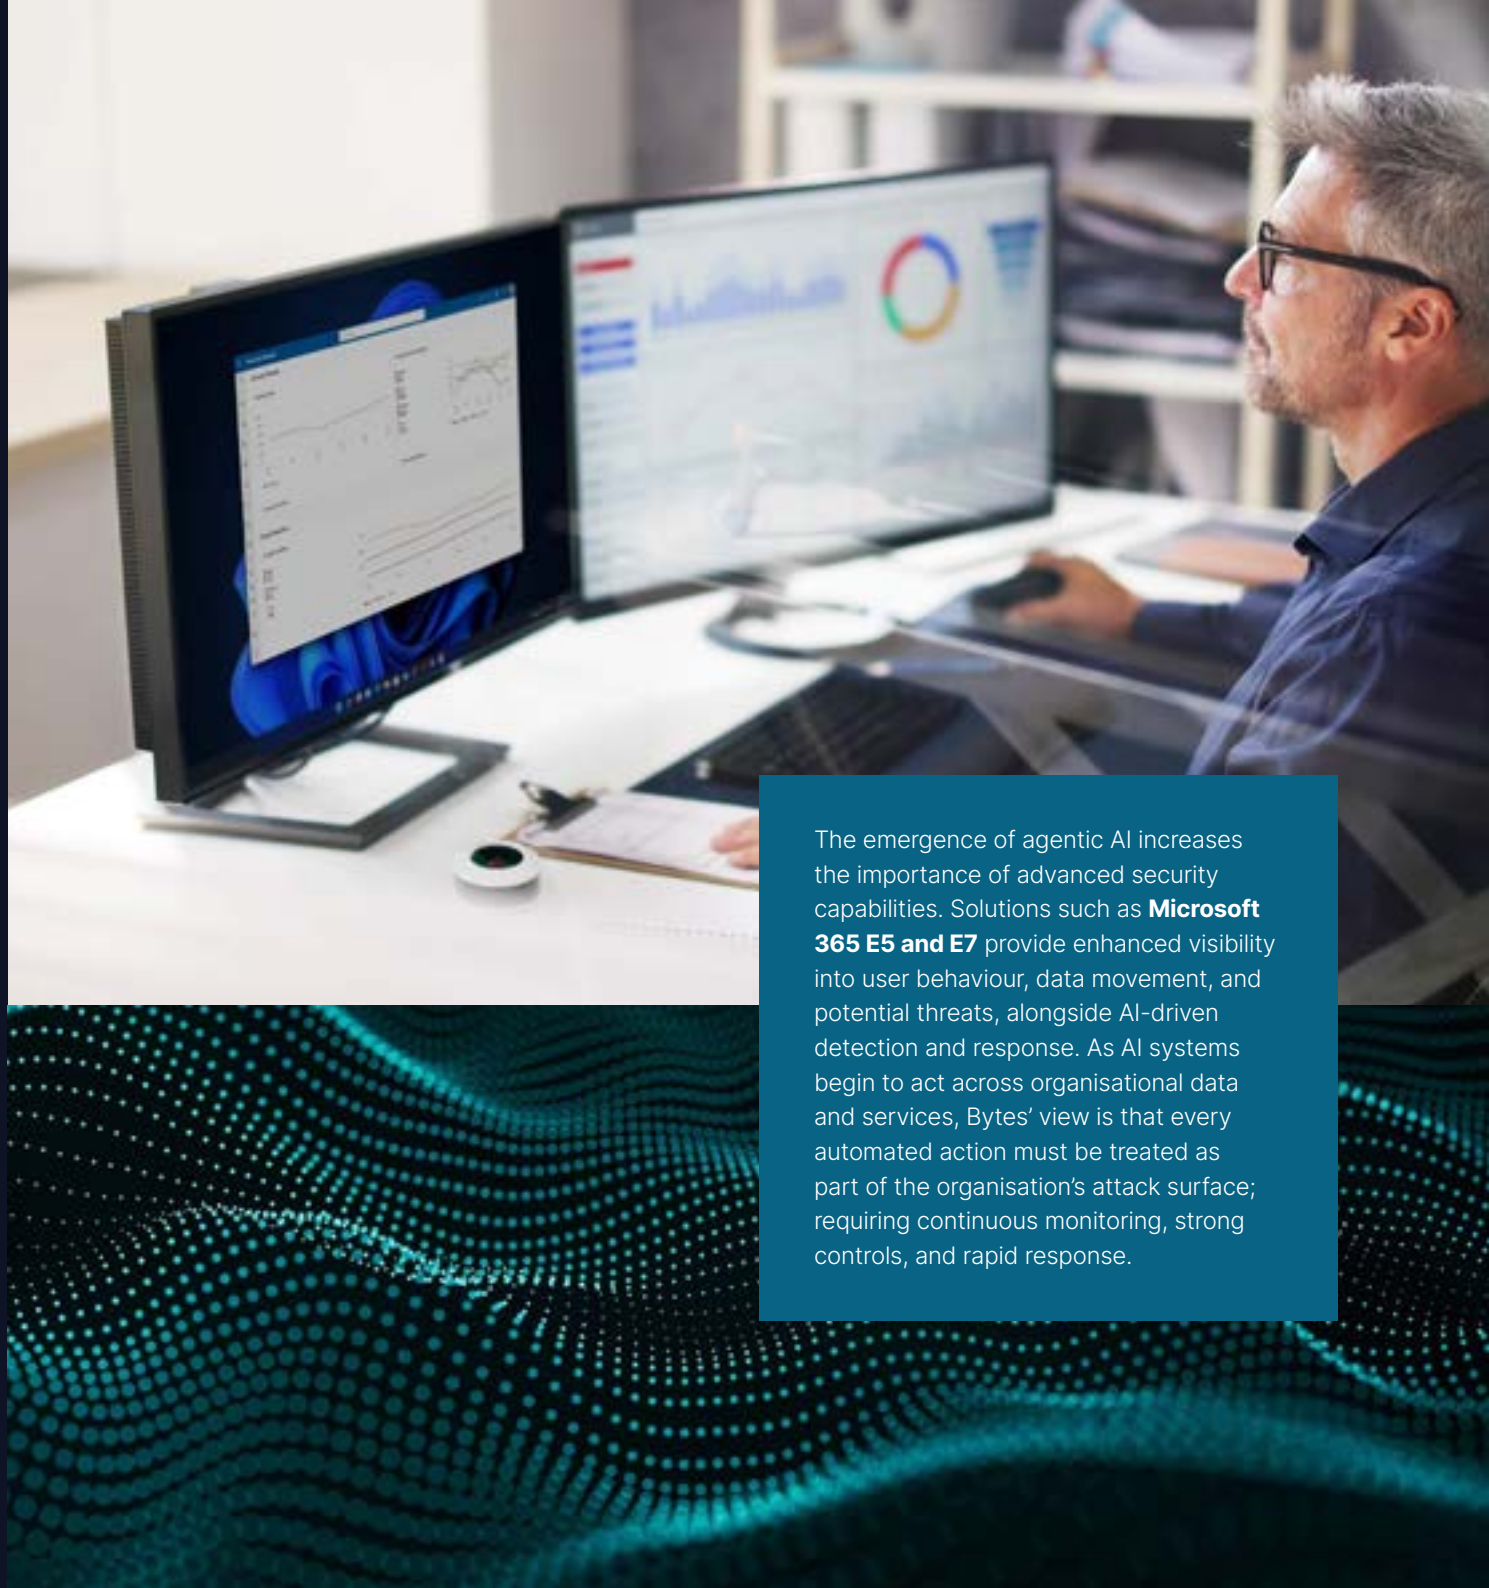

The emergence of agentic AI increases the importance of advanced security capabilities. Solutions such as **Microsoft 365 E5 and E7** provide enhanced visibility into user behaviour, data movement, and potential threats, alongside AI-driven detection and response. As AI systems begin to act across organisational data and services, Bytes' view is that every automated action must be treated as part of the organisation's attack surface; requiring continuous monitoring, strong controls, and rapid response.

Common failure patterns

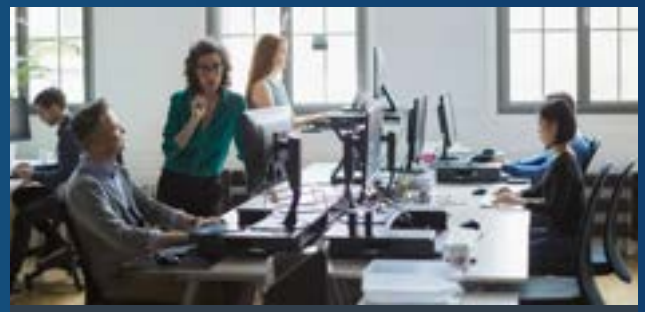

No visibility tools in place, leaving organisations blind to shadow AI.

Failing to make the connection between shadow AI and cyber risk.

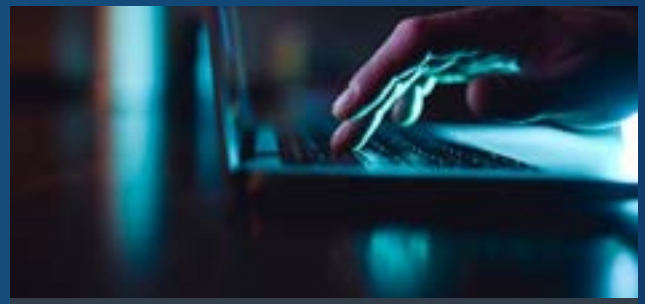

Reactive rather than proactive monitoring detects incidents only after they've occurred.

Data loss prevention policies are too broad (blocking legitimate work) or too narrow (missing risky behaviour).

Strategic outcomes

When organisations master this pillar, they have visibility into AI-related risks and can act before they escalate. Data leakage is detected and blocked automatically. The organisation demonstrates due diligence, meeting regulatory expectations for data protection, information governance and cyber resilience.

Operationalising the Framework

The five pillars set the strategic direction.
Implementation is where the real progress happens.

From Bytes' experience, organisations rarely tackle shadow AI all at once. Most begin by focusing on visibility, understanding the scale and nature of shadow AI before designing controls or deploying new tools. This discovery phase gives leadership the evidence they need to make informed decisions about where to focus first.

From there, organisations typically move towards putting guardrails in place. This is about establishing policies, enforcing technical controls, and communicating expectations to staff. The focus shouldn't be about punishment - it's about containing risk while the organisation prepares to offer better alternatives.

The longer-term goal is enabling safe, sanctioned AI that meets the same needs staff are currently fulfilling with consumer tools. When that's achieved, shadow AI doesn't need to be suppressed - it simply becomes unnecessary.

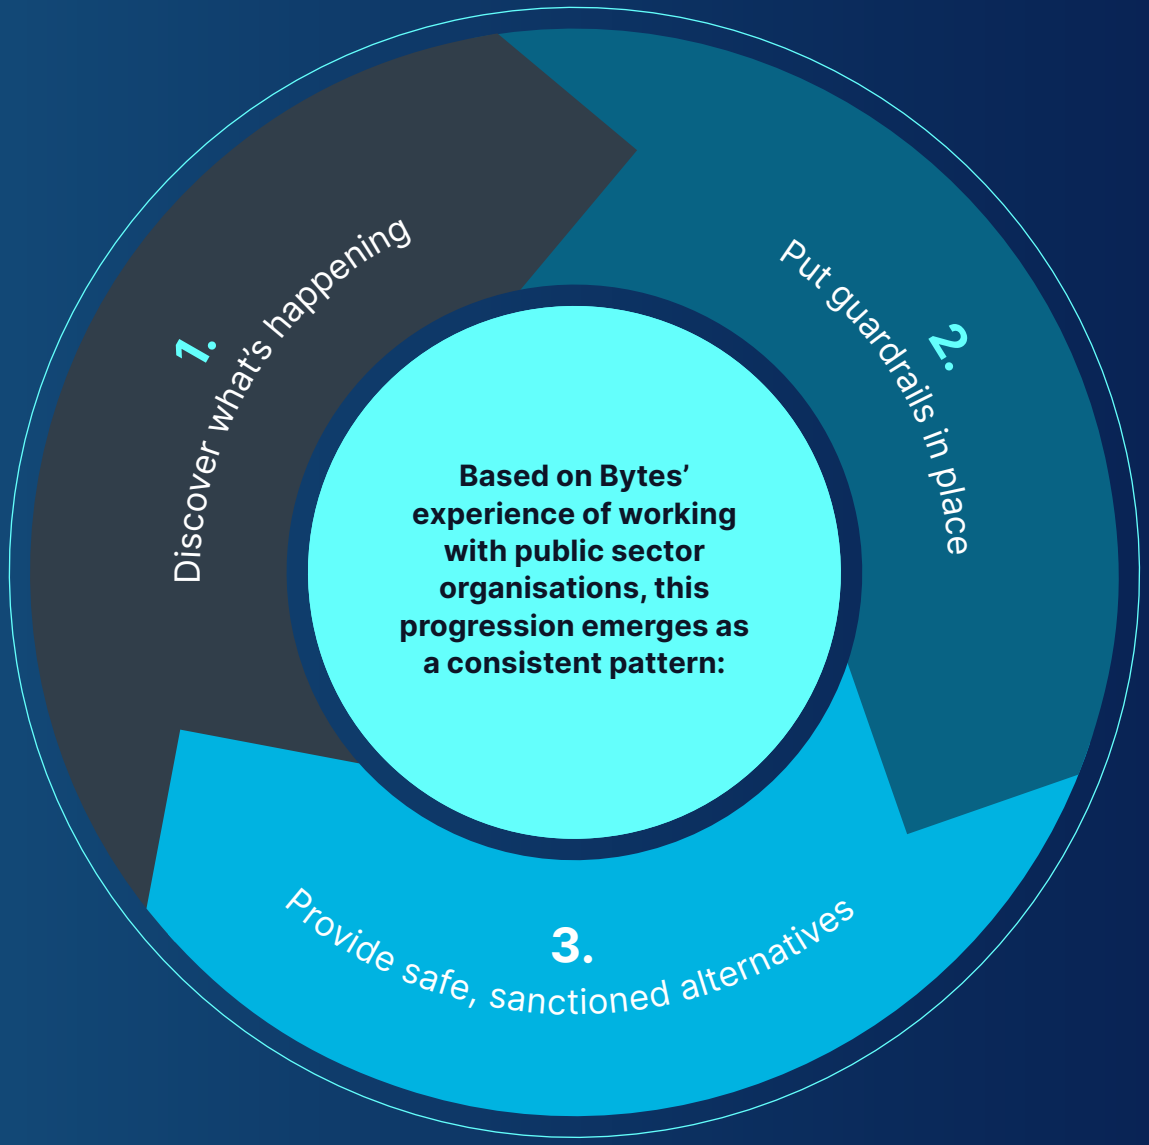

This three-phase approach represents how organisations typically execute against the Shadow AI 2026 framework in practice. The five pillars provide the strategic model, the 'what' and 'why' of addressing shadow AI. The three-phase approach is the execution pattern, the 'how' and 'when.'

The process isn't rigid or linear. Different organisations move at different speeds depending on their starting point, resources, and risk appetite. Some will need to revisit earlier stages as their AI landscape evolves. But the direction of travel is from invisible risk to managed, and value-driven AI adoption.

From Shadow AI to Strategic AI

With the right approach, shadow AI can be reframed from a risk to a strategic opportunity. It's an indication of appetite for AI. This willingness to innovate should be championed, not punished.

Bytes believes that suppressing shadow AI will only backfire. It's human nature that people will find workarounds. A blanket ban risks driving shadow AI underground. With the availability and familiarity of consumer tools, the risk is that staff will default to their personal devices for work use, multiplying ethical and legal issues.

Instead of suppressing shadow AI, organisations need to shift to strategic enablement. This reduces risk organically while rewarding those who are ready and willing to innovate. The goal is not to eliminate AI usage - it's to make it safe, visible and valuable.

From AI Assistants to AI Agents: A New Phase of Shadow AI Risk

AI in the workplace is entering a new phase. Where early adoption has centred on tools that support individuals (drafting, summarising and analysing) organisations are now beginning to deploy AI agents that can plan, act and automate tasks across multiple systems.

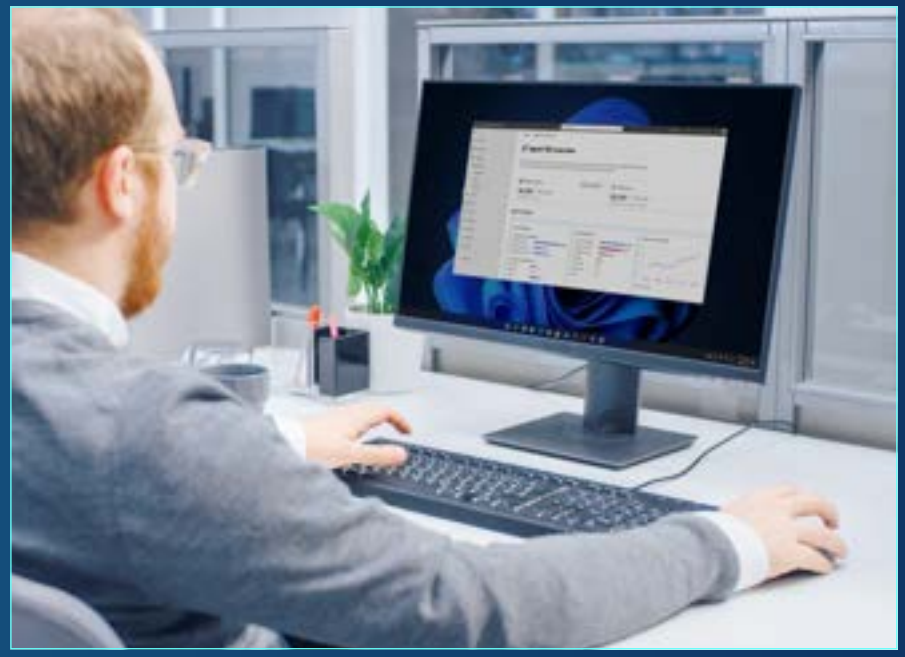

This creates clear opportunities to improve productivity and service delivery. However, it also fundamentally changes the risk profile.

Bytes believes that agentic AI has the potential to accelerate the shadow AI challenge. When AI systems can take action, not just generate content, any gaps in governance, data control or oversight are amplified at scale. What was previously a user-level risk becomes an organisational one.

Capabilities such as **Microsoft Agent 365** introduce a central control layer for managing AI agents across both Microsoft and third-party environments. This enables organisations to discover, govern and monitor agent activity; bringing visibility to what would otherwise become a new form of Shadow AI.

This shift places greater emphasis on enterprise-grade security. Solutions such as **Microsoft 365 E5 and E7** extend this control with advanced visibility and protection, including data loss prevention, insider risk management, audit, and AI-driven threat detection and response.

In an agent-driven environment, every AI action becomes part of the organisation's operational and security landscape. Leaders must ensure that AI is not only enabled, but visible, governed, and accountable. For the public sector, the implications are significant. Without the right foundations, today's shadow AI risk could evolve into autonomous, invisible decision-making at scale... beyond the reach of traditional controls.

Bytes' view is clear; the challenge is no longer just adopting AI safely, but operating AI systems that can act without losing control.

What Good Looks Like

Addressing shadow AI may sound like a technical challenge. Yet the benefits reach far beyond IT, transforming how leadership makes decisions, ways of working for staff, and ultimately, service delivery.

With sustained commitment, the signs of success are visible across the whole organisation:

Clear leadership ownership

AI strategy is owned by a senior leader, communicated regularly, and integrated into organisational objectives and strategies.

Confident, AI-literate workforce

Staff understand AI's capabilities and limitations, know how to use approved tools effectively, and recognise when to seek support.

Secure, approved AI embedded in workflows

Microsoft 365 Copilot is used daily for drafting, summarising, analysing, and automating routine tasks, all within a governed environment.

Public trust maintained

The organisation demonstrates responsible AI use, meets regulatory obligations, and can explain and defend its approach to citizens, regulatory bodies, and the media.

Dramatically reduced Shadow AI usage

Unsanctioned AI tools are rarely used because the sanctioned alternative is better. Where exceptions exist, they're visible, assessed, and managed.

Measurable productivity gains

Time saved is tracked, benefits are communicated, and AI adoption is tied to service improvements and operational efficiencies.

Strong governance without stifling innovation

Policies are clear, approval processes are streamlined, and usage is monitored without micromanagement.

Conclusion: A Leadership Imperative

Shadow AI is already happening. Across the public sector, staff are using consumer AI tools to draft reports, summarise intelligence and automate routine tasks. Much of this usage is invisible to IT teams - and to senior leaders. Yet ignoring the issue only increases the risk.

A human problem demands a human response. Leaders must start with understanding what's driving the behaviour. Shadow AI signals that staff are looking to work more efficiently and are open to innovation. Organisations risk demoralising staff by punishing early adopters who want to embrace AI. Excessive restrictions will only drive AI usage further underground, losing visibility and control entirely.

Bytes believes that strategic enablement is the only sustainable response. Organisations can regain control by providing better alternatives. This requires the right leadership and skills, data readiness, integrated AI platforms like Microsoft CoPilot, clear governance and robust cyber defences. The Shadow AI 2026 framework provides the roadmap to move organisations from reactive crisis management to proactive capability building.

Shadow AI can remain a hidden liability - or become the foundation of responsible, strategic AI in the public sector.

Ready to Move From Shadow AI to Strategic AI?

Most public sector organisations are already experiencing Shadow AI, even if they can't see it. Bytes can help you identify your real exposure and move quickly to safe, strategic AI adoption.

From supporting public sector clients from HMRC to NHS England, Bytes understands the public sector realities. In complex, pressured and highly-regulated environments, a focus on safety, scale, and sustainability is essential. It's all about turning risks into capabilities.

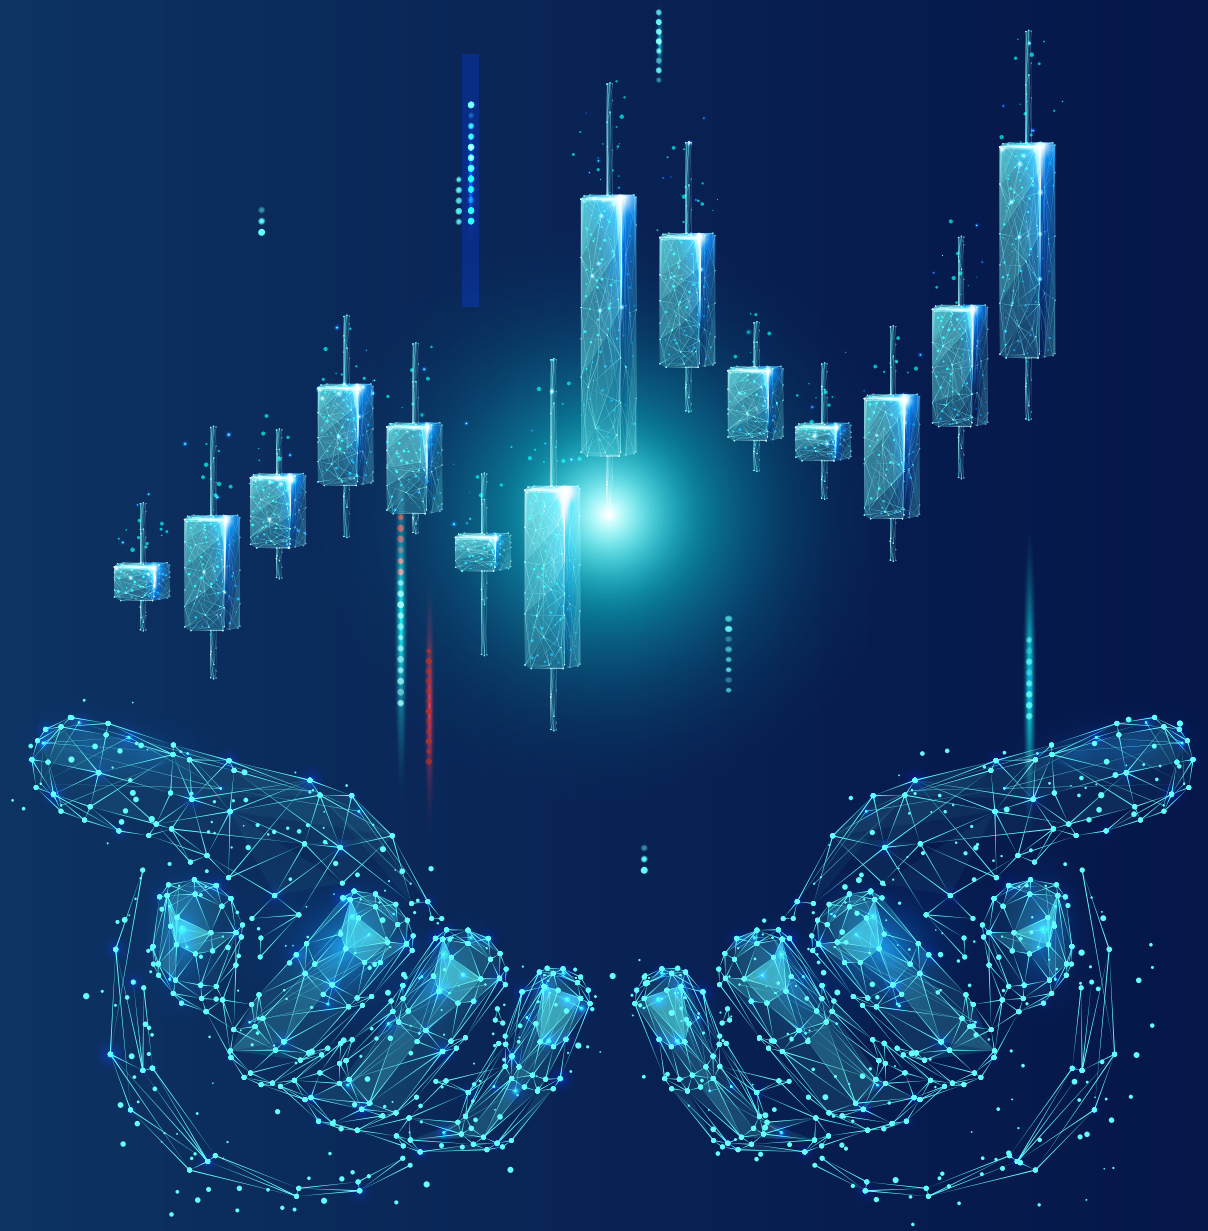

Bytes can help you through:

Shadow AI readiness assessments
to understand the scale and nature of risk across the organisation.

AI governance frameworks
that enable innovation without compromising security and oversight.

Prompt-a-thons and adoption programmes
to build capability and surface high-value use cases.

Leadership briefings
to build a clear understanding of the risks and how to address them.

Copilot readiness workshops
to ensure data governance and infrastructure foundations are in place before deployment.

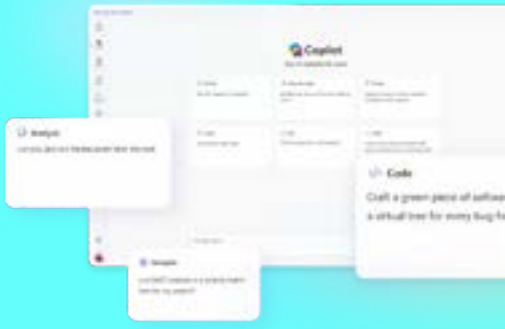

Executive workshops on AI operating models
to give leadership the clarity and confidence to drive AI adoption across the organisation.

Cross-department governance discussions
to ensure AI strategy is owned and understood across the organisation, not just within IT.

Bytes can support you in your journey from shadow AI to strategic AI. Email **tellmemore@bytes.co.uk** to discuss the next steps for your organisation.

References

1, 7, 12, 13
Microsoft. 'Rise in shadow AI tools raising security concerns for UK.'
ukstories.microsoft.com/features/rise-in-shadow-ai-tools-raising-security-concerns-for-uk/

2
Gov. uk. 'AI Skills for Life and Work Survey.'
www.gov.uk/government/publications/ai-skills-for-life-and-work-drivers-analysis/ai-skills-for-life-and-work-drivers-analysis--2

3
Gov.uk. 'AI Opportunities Action Plan.'
www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan

4
NHS England. 'The 10-Year Health Plan.'
www.england.nhs.uk/long-term-plan/

5
Ministry of Defence. 'The Defence AI Playbook.'
assets.publishing.service.gov.uk/media/65bb75fa21f73f0014e0ba51/Defence_AI_Playbook.pdf

6
Deloitte. 'Generative AI Consumer Trends.'
www.deloitte.co.uk/mediatelecomsbeyond/assets/pdf/deloitte-uk-generative-ai-dct-jun-24.pdf

8
The Telegraph. 'GPs use ChatGPT to help them treat patients.'
www.telegraph.co.uk/news/2024/09/18/gps-use-chatgpt-to-help-them-treat-patients-harvard-study-w/

9
Fortune. 'DWP staff banned from using ChatGPT.'
fortune.com/europe/2024/02/14/uk-government-dwp-social-security-banned-staff-using-chatgpt-microsoft-copilot/

10
Business Insider. 'OpenAI removes ChatGPT feature over search engine privacy concerns.'
www.businessinsider.com/openai-removes-chatgpt-feature-over-search-engine-privacy-concerns-2025-7?

11
The Guardian Australia. 'A ban ordered after child protection worker used ChatGPT in Victorian court case.'
www.theguardian.com/australia-news/2024/sep/26/victoria-child-protection-chat-gpt-ban-ovic-report-ntwnfb?

14
SOCITM. 'Shadow AI in the public sector.'
socitm.net/resource-hub/blog/shadow-ai-in-the-public-sector-innovation-without-oversight/

DELIVERED IN PARTNERSHIP

BYTES

At Bytes we are dedicated to helping organisations harness technology to achieve more. Our mission is to enable customers to modernise, innovate, and operate securely by providing the expertise, services, and solutions needed to succeed in a rapidly evolving digital landscape.

Bytes is a leading provider of world-class IT solutions, supporting organisations across the UK and Ireland with cloud, security, licensing, software asset management, virtualisation, storage, and managed services. Through deep technical expertise and close partnerships with leading technology vendors, we help customers transform their digital environments and maximise the value of their technology investments.

Our Microsoft Practice is one of the most established in the UK, built on more than 30 years of partnership with Microsoft. Bytes holds all six Microsoft Solutions Partner designations and was recently named a Microsoft Inner Circle Partner for AI Business Solutions, recognising our ability to deliver transformative solutions across Microsoft Cloud, Dynamics 365, and AI-powered innovation.

www.bytes.co.uk

Microsoft

At Microsoft we are dedicated to advancing human and organisational achievement. Our mission is to empower every person and every organisation on the planet to achieve more and by collaborating with policymakers around the world in addressing online security challenges, Microsoft supports global efforts to make the future of computing more secure.

Microsoft's solutions help address security issues and use AI and automation to detect and stop attacks automatically without human intervention. Get a holistic view into your environment and eliminate gaps in coverage with comprehensive cyber security solutions that work together and with your ecosystem to safeguard your identities, endpoints, apps, and clouds.

Today's world is more connected than ever before. Microsoft enables productivity and innovation by giving people the right solutions and processes to allow governments to take advantage of technology to improve how they communicate and deliver services without increasing the risk of attack.

www.microsoft.com

GOV.NEWS

At GovNews we are dedicated to strengthening collaboration between the public and private sectors. Our mission is to support innovation, knowledge sharing, and meaningful partnerships that help organisations across government improve services and deliver positive outcomes for citizens across the UK.

GovNews is a trusted Public Sector news and engagement platform that connects public servants with the ideas, technologies, and partners shaping the future of government. Through our journalism, thought leadership, digital media, and events, we provide a space where leaders can share insight, explore solutions to complex challenges, and highlight innovation taking place across the Public Sector.

Our latest innovation, GovNews Community, is an exclusive digital platform designed for UK Public Sector professionals. The Community enables members to access expert insight, watch video content, participate in discussions, and connect with peers across government. By bringing together content, collaboration, and networking in one place, GovNews Community helps public servants share knowledge, stay informed, and drive meaningful progress across the sector.

www.gov.news