Bring Shadow AI under control before it controls you

NSG helps UK SMB and mid-market businesses identify unmanaged AI usage, assess risk, define approved tools, strengthen data governance and build practical AI workflows that people can actually use.
We help you move from hidden experimentation to controlled AI adoption across productivity, GTM, revenue operations and commercial workflows

SHADOW AI IS NOT A STAFF PROBLEM. IT IS A CONTROL FAILURE

Shadow AI is already inside most businesses.

Not formally. Not neatly. Not with a board-approved roadmap and a tasteful slide deck. More likely through employees using public AI tools, browser extensions, meeting assistants, writing aids, plug-ins and file summarisation apps because they are trying to get work done.

That is the uncomfortable truth.

Shadow AI does not usually start because employees are reckless. It starts because people are under pressure. They need to write faster, summarise quicker, respond better, research more efficiently, prepare reports, clean up notes, draft emails, review documents and deal with the daily operational mess that businesses politely call “ways of working”..

AI helps immediately.

So people use it.

The problem is not that employees want productivity. That part is rational. The problem is that many organisations have not created a safe, approved and practical route for AI use.

When the official route is slow, unclear, restrictive or absent, staff create their own.

That is Shadow AI.

And it is not a staff problem.

It is a control failure.

What Shadow AI actually means

Shadow AI is the use of AI tools, apps or features outside the organisation's approved technology stack, governance model, security controls or management oversight.

That includes obvious examples, such as employees pasting content into public AI chat tools. It also includes less obvious examples, such as AI meeting note-takers added to calls without approval, browser extensions connected to email, writing tools used for proposals, AI plug-ins attached to productivity platforms, or applications granted access to Microsoft 365 data without proper review.

The risk often starts small.

Someone uses AI to rewrite an email. Someone asks a tool to summarise a meeting. Someone uploads a document to get the key points. Someone uses a browser extension to improve a proposal. Someone asks AI to turn rough notes into a client update.

Individually, these actions can seem harmless.
The issue appears when sensitive data enters the wrong system.

Client information. Pricing. HR records. Internal reports. Board material. Sales notes. Contract terms.

Customer complaints. Case details. Medical, education, finance or legal data. Internal documents that should never leave the organisation's control.

Once that information has been entered into an unapproved tool, the business may have limited clarity on where it went, how it is stored, whether it is retained, who can access it, whether it is used for training, what jurisdiction applies and whether the supplier's terms meet the organisation's obligations. That is not a theoretical risk.

That is what happens when productivity tools outrun governance.

Why Shadow AI is rising

Shadow AI is rising because AI solves immediate problems faster than most organisations can respond.

People are being measured on speed, output and quality. They are expected to do more with fewer resources. They are drowning in messages, meetings, documents, reporting tasks and admin. Then along comes a tool that can draft, summarise, rewrite, classify and explain in seconds.

It is not hard to see why people use it.

The bigger issue is leadership delay.

If the business does not provide approved AI tools, usable guidance and role-specific examples, people will not wait politely while a committee schedules its third exploratory session. They will find a tool that works and use it.

This is especially common in busy operational environments, multi-site businesses and teams where people are under pressure to keep work moving.

That creates a simple lesson:

Shadow AI grows where official AI adoption is too slow, too vague or too impractical.

You cannot fix that with a stern memo.

You fix it by giving people a better route.

The wrong response: ban it and hope

Some leaders respond to Shadow AI by trying to ban AI use completely.

This rarely works.

A blanket ban may reduce visible usage, but it does not necessarily reduce actual usage. It simply pushes behaviour underground. Staff become less likely to disclose what they are using, less likely to ask for guidance and less likely to raise concerns.

That is worse.

The business loses visibility. IT loses control. Security loses the chance to manage risk. Leadership loses the ability to understand where AI is already creating value.

A ban may be appropriate for specific tools, data types or high-risk use cases. But banning AI as a general position is usually a sign that the organisation has no operating model for adoption.

The better response is controlled enablement. That means:

approved tools, plain-English policies, data handling rules, role-based guidance, monitoring, training, escalation routes and practical examples of safe AI use.

People should not have to guess.

The approved route must be easier than the unofficial one.

That is the core principle.

The real risks behind Shadow AI

Shadow AI creates several practical business risks.

The first is data exposure. Prompts and outputs can contain sensitive information. If that information enters an unapproved tool, the organisation may lose control over retention, processing, access and onward use.

The second is app permission risk. Many AI tools ask for access to files, calendars, mailboxes, calls, messages or cloud storage. A user may grant permissions without understanding the impact. Once that happens, the risk is not limited to what they type into the tool. The tool may be able to access wider business data.

The third is identity and access risk. If permissions across Microsoft 365, SharePoint, Teams, OneDrive or other systems are poorly managed, AI tools can expose overshared data. The AI does not create the permission failure. It reveals it at speed.

The fourth is audit risk. If an AI tool influences a decision, recommendation, communication or customer output, the business may need to explain how it was used. That is difficult when staff are using tools with no central logging, retention, review trail or reporting.

The fifth is supplier risk. Some AI tools may look harmless but have weak contractual terms, poor data controls, unclear sub-processor arrangements, unsuitable retention policies or limited breach reporting.

The sixth is commercial risk. AI-generated customer emails, proposals, campaigns or reports may contain inaccurate claims, weak judgement or unsupported statements. Because AI can make poor content look polished, it can create a false sense of quality.

That last point matters.

A bad human draft often looks bad.

A bad AI draft can look professional enough to escape challenge.

Apparently presentation quality remains one of humanity's favourite traps.

Shadow AI and UK data expectations

For UK businesses, Shadow AI needs to be considered through the lens of UK GDPR, data protection obligations, supplier due diligence and secure-by-design principles.

The Data Use and Access Act 2025 received Royal Assent on 19 June 2025 and updates areas of the UK's data and digital information framework.

Government guidance states that changes are being phased in, while ICO guidance explains that the Act changes some data protection laws affecting organisations using personal information.

The practical point is not that every AI use case requires legal panic.

The practical point is that businesses must know what personal data is being used, where it is being processed, who can access it, what lawful basis applies, what supplier terms say, and what safeguards are in place.

Shadow AI makes that difficult because the business may not even know which tools are being used.

That is the issue.

You cannot govern what you cannot see.

The Microsoft 365 misconception

Many organisations assume that because they are Microsoft 365 businesses, their AI risk is automatically controlled.

That is not true.

A Microsoft-first environment can provide strong foundations for AI governance, especially through identity, access control, data classification, sensitivity labels, Defender, Purview, SharePoint governance and Copilot readiness work.

But the platform does not fix poor configuration by itself.

If SharePoint permissions are messy, AI may surface information too widely.

If users can grant high-impact app permissions without proper review, third-party tools may access sensitive data.

If data is poorly labelled, it is harder to apply meaningful controls.

If employees do not understand what they can and cannot enter into AI tools, mistakes will happen.

If Copilot or other AI assistants are rolled out without readiness work, the organisation may simply accelerate existing information governance problems. Microsoft controls can help.

They do not replace leadership, governance or data discipline.

This distinction matters because tools do not save a business from its own lack of ownership. Cruel, but efficient.

Shadow AI and Commercial Transformation

Shadow AI is a Commercial Transformation issue because it exposes how work is really happening inside the business.

It shows where people are under pressure. It shows where processes are too slow. It shows where documentation is weak. It shows where knowledge is trapped in people's heads. It shows where teams need faster drafting, summarisation, research, analysis or reporting support.

In that sense, Shadow AI is not only a risk signal. It is also an operating signal.

If employees are using AI to summarise meetings, the business may have a meeting and action-tracking problem.

If salespeople are using AI to write follow-up emails, the business may need better templates, messaging and CRM workflow.

If managers are using AI to create reports, the business may have weak reporting structure or poor data visibility.

If teams are uploading documents for summarisation, the business may need better knowledge management and document intelligence.

If staff are using AI to rewrite customer communication, the business may need stronger service communication standards.
Shadow AI tells leaders where friction already exists.

The answer is not to ignore that signal.

The answer is to turn it into a controlled improvement programme.

Shadow AI and AI, Automation and Orchestrational Intelligence

For NSG, Shadow AI sits directly inside AI, Automation and Orchestrational Intelligence.

The opportunity is to move from unmanaged tool use to designed AI capability.

That means identifying where AI is already being used, assessing the risk, understanding the business reason behind the behaviour, and then designing safe, approved workflows that solve the same problem properly.

For example:
If employees use public AI to summarise internal documents, the controlled route may be a secure document intelligence workflow.

If teams use AI meeting tools informally, the controlled route may be an approved meeting capture, summary and action workflow with clear consent and retention rules.

If sales teams use AI to draft outreach, the controlled route may be approved prompt structures, messaging guardrails and review workflows.

If managers use AI for reporting narratives, the controlled route may be governed data sources, standardised report templates and human review.

If operations teams use AI to interpret policies, the controlled route may be a secure internal knowledge assistant with approved source material.

The point is not to remove AI.

The point is to remove unmanaged AI.

Shadow AI and Revenue Operational Control

Revenue functions are especially exposed to Shadow AI risk.

Sales, marketing, customer success and account management teams deal with commercially sensitive information every day. That includes pipeline data, pricing, proposals, customer pain points, negotiation details, qualification notes, call recordings, forecasts, account plans, competitor references and renewal risks.

If that information is pasted into unapproved AI tools, the business may lose control over sensitive commercial intelligence.

The risk is not limited to data protection.
It also affects revenue integrity.

AI-generated follow-ups may misrepresent what was agreed.

Proposal drafts may include unsupported claims.
Forecast summaries may sound stronger than the evidence supports.

Account research may rely on inaccurate sources.
Outreach personalisation may become intrusive or wrong.

CRM notes may be summarised without preserving important context.

Deal risk may be missed because the underlying data was weak.

Revenue control depends on accurate information, disciplined process and trustworthy handoffs.
Shadow AI can undermine all three.

The better approach is to create approved AI workflows for revenue teams:
call summary templates, qualification review prompts, proposal drafting frameworks, approved messaging libraries, CRM completeness checks, follow-up standards, forecast commentary support and deal risk review prompts.

That is how AI strengthens revenue control.

Not by letting every salesperson improvise with whatever tool happens to be open in their browser.

Shadow AI and Growth GTM

Shadow AI also affects Growth GTM.

Marketing and GTM teams are under constant pressure to produce content, campaigns, messaging, insight, research, segmentation and outreach at speed. AI is obviously useful here.

But unmanaged AI use can create serious problems. Generic content. Weak differentiation. Unsupported claims. Poor targeting. Incorrect market assumptions. Over-personalised outreach. Misuse of prospect data. Brand inconsistency. Duplicate content. Poor source control. Campaign volume without commercial relevance.

AI can make GTM faster.

It can also make bad GTM faster.

That is the issue.

The aim should not be to use AI to produce more noise. The aim should be to use AI to improve market focus, buyer relevance, message quality and campaign discipline.

That requires controlled source material, approved messaging architecture, clear claim boundaries, proper review, and a direct link between campaign output and sales follow-up.

Shadow AI breaks that link.

Governed AI strengthens it.

How leaders should bring Shadow AI under control

The practical answer is not complicated, but it does require discipline.

First, discover what is already happening.
Do not assume staff are waiting for permission. They are probably not. Run an AI usage review. Ask teams what they are using, what they are using it for, what problems it helps solve and what data they enter. Make it safe for people to be honest. If the review feels like a disciplinary trap, people will hide the truth, because apparently self-preservation remains popular.

Second, classify the use cases.
Separate low-risk drafting and productivity use from higher-risk use involving personal data, customer information, contracts, financial information, HR records, regulated content, automated decisions or external communication. Third, define approved tools. Give employees clear options. If the business wants people to stop using random public tools, it must provide approved alternatives that are easy to access and good enough to use.

Fourth, create plain-English rules.
Policies should explain what people can do, what they cannot do, what data must never be entered, what needs review and where to ask for help. Role-based examples matter more than abstract principles.

Fifth, fix permissions and data access.
Before scaling AI across Microsoft 365 or other internal systems, review access rights, oversharing, sensitive data exposure, retention, labelling and document governance. AI will surface what your permission model already allows.

Sixth, monitor without making work impossible.
Security monitoring should identify risky tools, high-risk access patterns and unusual activity. But it should not create so much friction that people go back to hidden behaviour.

Seventh, build safe workflows.
Turn common Shadow AI behaviours into approved workflows. If people are using AI for meeting notes, create a controlled meeting summary process. If they are using AI for proposals, create an approved proposal drafting workflow. If they are using AI for customer responses, create review rules and tone standards.

Eighth, train managers, not just users.
Managers need to understand how AI changes work quality, review standards, accountability and performance expectations. Otherwise, employees will use AI while managers judge output as if nothing has changed

Ninth, measure value and risk.
Track whether approved AI routes reduce admin, improve consistency, reduce errors, speed up workflows or strengthen control. Also track incidents, risky behaviours and exception cases.

Tenth, keep governance live.
AI tools, risks and use cases will change. A policy written once and ignored for two years is not governance. It is decorative paperwork.

A practical Shadow AI control framework

A useful Shadow AI control model has six layers.

1. Discovery
Identify which AI tools are being used, by whom, for what purpose and with what data.
2. Risk classification
Separate low-risk productivity use from high-risk data, customer, employee, financial, legal or regulated use.
3. Approved routes
Provide safe AI tools and workflows that solve the same productivity problems without exposing the business.
4. Data and access control
Review permissions, sensitivity labels, retention, sharing settings, supplier terms and system integrations.
5. Human oversight
Define which outputs need review, who approves them and when AI must escalate to a person.
6. Monitoring and improvement
Use monitoring, reporting, feedback and governance reviews to keep the model current.

This is the sensible middle ground.

Not panic.

Not free-for-all.

Controlled enablement

What businesses should not allow

Some AI use should be clearly prohibited unless a specific, reviewed and approved workflow exists.

Employees should not enter personal data, customer records, HR details, financial data, board papers, contracts, source code, credentials, confidential client material or commercially sensitive documents into unapproved AI tools.

They should not connect AI apps to business systems without approval.

They should not use AI meeting assistants on external calls without clear rules around consent, notification and data handling.

They should not send AI-generated external advice without review.

They should not rely on AI outputs for legal, HR, financial, medical, regulatory or contractual decisions without appropriate human oversight.
They should not use AI-generated claims in GTM or sales material unless those claims are evidenced and approved.

This needs to be explicit.

People cannot follow rules that leadership has not bothered to define.

NSG view: Shadow AI is a warning light, not just a risk

At Noodle Spark Group, our view is direct.

Shadow AI is not just a security issue.

It is a warning light showing that the business has unmet productivity demand, unclear governance, weak enablement or poor operating design.

The wrong response is to clamp down blindly.

The right response is to understand why people are using AI unofficially, then build a safer, better and more commercially useful route.

That means connecting AI adoption to real business work:
sales follow-up, proposal creation, customer communication, reporting, meeting actions, knowledge management, operational workflows, revenue control and GTM execution.

It also means applying proper controls:
approved tools, data rules, permissions, monitoring, human review, supplier due diligence and clear accountability.

Shadow AI becomes dangerous when leadership ignores it.

It becomes useful when leadership treats it as evidence.

Evidence of where work is too slow.

Evidence of where people need better tools.

Evidence of where governance is missing.

Evidence of where AI could create value if it was brought under control.

That is the practical path.

Not banning AI.

Not pretending everything is fine.

Not waiting for the next compliance scare to create urgency.

Bring the usage into the open. Govern it properly. Replace unsafe workarounds with approved workflows. Use the insight to improve how the business operates.

That is how Shadow AI turns from unmanaged risk into controlled commercial capability.