wavenet

Microsoft Solutions Partner

Explore the possible

How organisations can adopt AI without losing control

A practical perspective on leadership, data and safe adoption

Contents

Introduction:

AI is already part of everyday work

If you work in a modern organisation, AI is already in the building. Maybe not officially or consistently, but it's undoubtedly there.

Microsoft has reported that 75 percent of knowledge workers now use AI at work, and 78 percent of AI users are actually bringing their own tools into work because they're tired of waiting for businesses to catch up with the curve. This reflects what businesses are telling us week in, week out in 2026. Staff are under mounting pressure, inboxes continue to grow, and there's always more work than time. Teams are looking desperately for ways to draft quicker, summarise faster, and reach decisions earlier.

The problem definitely isn't curiosity at this point, it's what happens when experimentation spreads faster than guidance, training, and controls. That's when you start seeing risky workarounds, inconsistent output, and data being shared in risky places it doesn't belong in.

We help organisations stay innovative, connected and protected, and we've been doing that since 2000. We run a CREST and CHECK accredited cyber security division and a 24x7 UK security operations centre, so we tend to look at this through a safety and real-world lens. These are challenges we're seeing every week, and conversation's we're having day in day out.

This eBook is here to help you explore what's possible through embracing AI, give you a solid understanding of common misconceptions, and the risks of overlooking its potential.

Shadow AI:

What's commonly happening in organisations now

Before we talk about shadow AI, it's worth taking a step back and looking at what every day work looks like right now.

Most organisations are already full of people experimenting quietly. Someone pastes meeting notes into an AI tool to get a quick summary, a worker rewrites a tricky email, so it sounds clearer, while another asks ChatGPT to help structure a proposal and plan a presentation.

None of this feels particularly controversial, does it? In fact, it often feels helpful and sensible. People are under pressure to do more with less time, and AI tools promise quick answers and less admin.

The problem is that this AI experimentation usually starts before there's any formal guidance. There's no agreement on which tools are approved, what data is safe to share, or how usage should be monitored. So, people do what people always do - they use what's easiest, fastest and already familiar.

What that looks like:

- An employee uses a public AI tool because it's quicker than logging an IT request or waiting for guidance
- A team starts relying on one particular tool while another team uses something completely different
- Someone copies information into a prompt without fully thinking about where that information goes next

At this stage, no one thinks they are doing anything risky, and there's certainly no bad intent. Most people would be surprised to learn there's a potential issue at all.

This is an example of shadow AI in action.

Shadow AI is simply the name we give to this behaviour once it becomes widespread and invisible. It's what happens when AI tools are used at work without IT or security teams knowing, approving, or therefore able to apply the appropriate security and governance controls. In many ways, it's the next evolution of shadow IT, but it moves faster and is harder to spot.

Industry research shows just how common this has become. Gartner predicts that by 2030, more than 40 percent of organisations will suffer security or compliance incidents linked to the use of unauthorised AI tools. In the same survey, it was revealed that 69 percent of cyber security leaders either suspect or have evidence that employees are using public AI tools at work. The key thing to understand is that Shadow AI isn't usually driven by recklessness, but by pressure. People want answers faster, quicker drafts, and less time buried in admin.

Once you look at it through that lens, the behaviour makes sense.

What begins as helpful experimentation turns into something widespread that organisations can't see, measure, or control. Different tools are used for similar tasks. Data is shared inconsistently and unintentionally. IT and security teams are left reacting rather than guiding.

This is also why the instinct to simply block AI tools rarely works. People don't stop exploring. They just move elsewhere, and visibility gets worse rather than better as it becomes deliberately "under the radar".

Shadow AI isn't a failure of employees. It's a signal that curiosity and need have moved faster than guidance. And once you see it that way, the focus shifts. The goal is no longer to stop experimentation, but to make it safer, more visible, and easier to manage.

Why blocking AI creates more risk, not less

When organisations first become aware of uncontrolled AI usage, the natural reaction is often to shut it down. From a security or compliance perspective, this makes sense. If a tool introduces risk, remove the tool.

In reality, blocking AI without providing an alternative can create more problems than it solves.

The simple truth is that you can't secure what you can't see, and you can't govern what you're not aware of. When tools are blocked outright, people rarely stop trying to use AI. Instead, they look for workarounds. They switch to personal devices, personal accounts, or tools that sit outside managed environments.

Research into AI use at work reinforces this pattern. Microsoft has found that **more than half of employees** would not tell their manager they were using AI to complete work tasks. That figure alone should give leaders pause. When people feel unsure, unsupported, or restricted, they don't speak up. They go quiet – and do it anyway.

This is one of the hidden pitfalls of banning AI tools. It pushes activity underground.

Once that happens, several things follow:

- **IT and security teams lose visibility**
  Usage moves beyond corporate identity systems and audit trails, making it far harder to understand what data is being shared or where exposure exists.

- **Risk becomes fragmented and inconsistent**
  Different teams find different tools, each with different safeguards, terms and behaviours, creating a sprawling and unmanaged risk surface.

- **Employees stop asking questions**
  Curiosity doesn't disappear, it just becomes harder for leaders to guide it.

- **Organisations fall behind quietly**
  While internal debate continues, competitors who provide safe, approved routes start learning faster, building experience earlier and responding more effectively to market opportunities.

Blocking also creates a cultural problem. It frames AI as something dangerous that must be avoided, rather than a capability that needs to be handled responsibly. Over time, this makes it harder to introduce structured adoption later because trust has already been damaged.

This is why many organisations are now shifting their approach. Instead of 'ban and hope', they are moving towards 'sanction and guide'.

The aim is to give people a safer route to explore within systems that are already managed. Environments where identity, access controls, permissions and audit trails already exist.

When organisations do this well, a few important things change:

AI use becomes visible rather than hidden

Data risk can be assessed and reduced systematically

Teams start learning together rather than in silos

Leaders regain the ability to guide adoption instead of reacting to it

This approach doesn't remove risk entirely, but it does something far more important. It brings exploration back into the open, where it can be supported, shaped and managed properly.

When productivity improvements become a business problem

Most organisations don't question whether AI can be useful. Their challenge is how to introduce it in a way that benefits the whole business, not just individuals.

When AI adoption happens informally, without shared ownership or direction, productivity gains tend to show up unevenly. Some teams move quickly, others hold back, and ways of working begin to diverge.

Over time, a few patterns emerge.

- **Different teams produce different outputs with different quality and tone**
  Without shared expectations or examples, teams develop their own habits. What feels efficient for one group can feel inconsistent or confusing when viewed across the organisation.
- **People rely on different data sources**
  Answers and drafts may be technically accurate in isolation, but they are not always aligned. This can create subtle contradictions that only surface later.
- **Knowledge stays trapped in silos**
  Prompts that work well, lessons learned, and better ways of working rarely travel beyond the individual or team that discovered them.
- **Leaders lose clarity**
  When usage is scattered across tools and accounts, it becomes difficult to understand where AI is creating value, where it's creating duplication, and where it might be introducing new risk.

None of this means an organisation is doing anything wrong, it just reflects what happens when a new capability grows faster than it can be safely implemented.

If adoption is driven primarily by individuals, the result is usually individual efficiency rather than organisational effectiveness. People feel more productive, but processes don't improve in a consistent or repeatable way.

This is the point where productivity gains can quietly turn into a business problem. Not because AI has failed, but because it's being used without a shared frame of reference.

The reassuring part is that this is solvable.

With the right guidance, clear standards, and a considered approach to rollout, organisations can turn these early patterns into something far more coordinated, that builds confidence in AI as adoption grows.

The data exposure organisations underestimate

Most organisations don't get caught out by a major breach. They get caught out by small, everyday behaviours that create exposure over time.

What goes wrong most often, and what to check first

Here are the data exposure patterns that come up repeatedly, and the practical checks that reduce the risk:

Over-permissioned access

If your internal files and folders are too widely accessible, AI tools that respect permissions will simply make more content visible to more people. The fix isn't 'avoid AI', it's to align access permissions with how your organisation works.

Copy and paste of confidential information

Directly entering prompts into public tools is a common leakage route, especially when data may be retained or used beyond your control. That's why acceptable use rules need to be applied centrally, and clear and repeated guidelines issued, not hidden in a policy folder.

Personal accounts and unmanaged logins

When people use personal accounts, IT can't track and audit activity. Encouraging approved tools with single sign-on (SSO) improves visibility and control.

Unapproved extensions and add-ons

Some extensions can read content and access data across web apps, which can bypass traditional perimeter controls and expose information.

Lack of clear boundaries

Most people are not trying to take risks. They just don't know where the line is. This is why education and guidance matter as much as the tooling.

If you only take one action from this section, focus on visibility first. If you can't see what tools are being used and how, you can't reduce risk in a meaningful way.

Why AI initiatives stall and why leadership is the real unlock

All organisations start their AI journey with good intent. They run a pilot, test a few use cases and early results look promising. But then things slow down or stop altogether.

This usually happens for the same reasons:

- There's no clear owner across the business, so responsibility is fragmented
- Success is hard to define, so progress is difficult to measure
- The focus drifts towards licences and tools, rather than readiness and ways of working
- Risk feels present, but there is no shared plan for how to manage it

Research reflects this gap between ambition and preparedness. In Microsoft and LinkedIn's Work Trend Index, 60 percent of leaders say their organisation lacks a clear vision and plan for implementing AI. In other words, many organisations want the benefits but are still unsure how to move forward with confidence, so the stall isn't a surprise.

When leadership feels uncertain, teams tend to experiment quietly on the side. When experimentation happens without shared direction, risk increases. And when risk becomes harder to understand, leadership often pulls back further. Over time, progress slows to a crawl.

What gets missed in this cycle is that AI adoption can't sit with one enthusiastic individual or one motivated team. When it does, you end up with pockets of progress and a lot of inconsistency. Leadership makes the difference because it's the only way to move from individual activity to shared outcomes.

Clear leadership creates:

Shared standards for how tools should be used and what good looks like

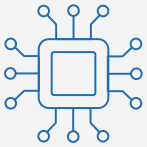

Agreement on where the value should show up first, so effort is focused

Consistent training and support, rather than ad hoc learning

Clear boundaries around data, compliance, and acceptable use

Without this direction, AI remains something people use on their own terms. Helpful in places but disconnected from wider business goals. With it, AI becomes something teams learn together, refine together, and build confidence in over time.

The way forward isn't about rushing into full deployment or trying to predict every outcome upfront. It's about creating a structured starting point that gives leadership confidence and gives teams clear guidance.

That shift, from individual experimentation to shared ownership, is what allows AI initiatives to move beyond pilots and turn into meaningful progress.

Microsoft 365 Copilot Launchpad:

The foundation for safe AI exploration

Safe exploration isn't about slowing everything down or overengineering the early stages of AI adoption. It's more about putting just enough structure in place so learning is visible, progress is repeatable, and risk is understood.

At its simplest, safe exploration works when three things move together: people, data and governance. When any one of these is treated in isolation, problems start to creep in. Teams experiment without shared guidance. Data permissions are overlooked until something goes wrong. Leaders struggle to see and understand what's really happening across the organisation.

What consistently works better is a more joined-up view of readiness.

That means taking time to understand where AI is already being used and why. It means prioritising education and clarity over punishment or restriction. And it means looking honestly at how identity, access and data permissions are set up before usage scales. When organisations approach this as a learning exercise rather than a policing exercise, trust increases and visibility improves. People are more open about what they are trying. Leaders get a clearer picture of where value is emerging and where support is needed.

This is where many organisations realise they don't need more theory about AI. They need a practical starting point that turns uncertainty into a plan. That's the role Microsoft 365 Copilot Launchpad is designed to play.

Why Microsoft 365 Copilot Launchpad

A great place to start is through a structured exploration of Microsoft Copilot within a managed Microsoft 365 environment, without rushing straight into a full rollout. Rather than issuing licences in isolation, it brings together leadership alignment, technical readiness, training and real use cases within a defined timeframe.

Delivered across four structured engagements over 30 days, Microsoft 365 Copilot Launchpad is designed to answer the questions that tend to hold organisations back at this stage. Where do we start? Who owns this? What needs to be in place before we scale, and how do we know if this is actually working?

In practical terms, this includes:

Leadership alignment and governance

Early sessions focus on clarifying intent, agreeing on priorities, and establishing shared ownership so AI adoption doesn't rely on individual enthusiasm alone.

A realistic readiness check

This covers licensing, identity and access, collaboration tools, and where organisational data lives today. The aim is to surface gaps early, while they are still straightforward to address.

Structured enablement

A cross-functional group is trained on how to use Copilot effectively and responsibly, covering foundations, prompt skills, and awareness of areas such as information sensitivity and bias.

Use cases grounded in real work

Rather than generic examples, the focus is on how Copilot can support day-to-day tasks within specific roles and teams, with a clear view of what meaningful impact looks like.

Clear next steps

The output isn't just insight, but a plan. What should be scaled, what needs fixing first, and where longer-term value is most likely to sit.

Crucially, this approach keeps exploration within systems that are already familiar and governed, where identity, permissions and audit trails exist. That makes it far easier to understand usage, reduce risk and build confidence over time.

Safe exploration doesn't remove uncertainty entirely. But it replaces guesswork with learning, and hesitation with a sense of progress. For organisations looking to move beyond curiosity without overcommitting too early, that balance matters.

Building confidence for what comes next

Safe exploration of AI is what allows organisations to move forward with confidence.

Once there is clearer governance, better visibility, trained people, and shared agreement on where AI should be applied, decisions become easier to make, and progress feels deliberate rather than reactive. And confidence starts to replace hesitation.

At this point, organisations aren't guessing anymore. They have a grounded understanding of where value sits, what needs attention before scaling, and how AI shows up in real, day to day work. That clarity makes it possible to move beyond isolated experimentation and into more consistent adoption across teams, without feeling rushed or overcommitted.

The organisations that progress most effectively tend to take a phased approach. Rather than trying to do everything at once, they build confidence first, strengthen the foundations that matter most, and then expand capability over time. That might mean extending Copilot into more roles, refining data readiness, or introducing more advanced automation once teams are genuinely ready for it.

This measured progression is what prevents pilots from stalling. Instead of jumping straight to full deployment, organisations create a safe first step, learn quickly from real use, and scale deliberately based on what they now understand. Progress becomes something that builds, rather than something that stalls under the weight of uncertainty.

This is where having the right partner becomes important.

Guidance when it matters most

We work with organisations that want to explore AI responsibly within their existing Microsoft environments. Being a Microsoft solutions partner we have in-depth experience across security, connectivity, cloud and communications, our focus is on making sure AI works in a real, operational context, not just in theory.

Microsoft 365 Copilot Launchpad plays a vital role in that journey. It provides the structure needed to move from curiosity to informed decision-making, while keeping exploration grounded and manageable. Just as importantly, it sets organisations up for what comes next by building confidence, capability, and clarity, rather than forcing premature commitments.

The goal isn't to reach an artificial endpoint with AI adoption. It's to create the conditions where organisations can keep moving forward at the right pace, with clear direction and sensible guardrails in place.

If you're thinking about how AI could support your organisation but want to do it in a way that feels considered rather than rushed, the best place to start is a conversation. One that focuses on what's realistically possible, right now, in your organisation.

Sources and references cited in this guide:

- **Microsoft and LinkedIn**
  2024 Work Trend Index: The state of AI at work

- **Infosecurity Magazine**
  Gartner: 40% of firms to experience shadow AI incidents by 2030

- **Forbes**
  Workers don't want bosses knowing they use AI, even as they bring their own tools to work

Start exploring what's possible with AI for your business today.

Book a consultation >

enquiries@wavenet.co.uk

0333 234 0011

wavenet.co.uk

wavenet

Microsoft Solutions Partner