



AI IN THE FLOW OF HUMAN AMBITION

The IT Leader's Guide to Scaling AI with Confidence

How to build the foundation for secure AI adoption,
agent readiness and measurable business impact.

For IT and security leaders scaling AI across the enterprise – with Microsoft 365
as the foundation, and the flexibility to use the right model for every scenario.



Introduction



Most enterprises are no longer deciding whether to adopt AI. They're discovering they already have – in dozens of disconnected tools, bought by different teams and governed by no one in particular.



Independent enterprise telemetry finds that a majority of AI use inside organisations still runs through personal, unmanaged accounts, outside IT's visibility.

NETSKOPE, 2025


For IT and security leaders, that is the real challenge of this moment. Your people are already using AI to draft content, analyse information, automate repetitive work and speed up decisions. The mandate now is broader: turn that scattered momentum into a secure, governed, enterprise-wide capability that delivers measurable business value – without standing up yet another stack to run.



How can I help?

The cost of getting it wrong is climbing. Left unmanaged, independent adoption produces fragmented data, inconsistent governance, duplicated spend and blind spots in security – and it does so exactly as agents begin to take multi-step action across your systems. The risk is no longer a missed productivity gain; it's losing visibility and control at the moment both matter most.

The organisations pulling ahead share one move: instead of building a parallel AI stack, they extend the platform they already trust. This guide lays out how and where Microsoft 365 fits as that foundation.



Scaling AI takes more than adding tools.
It rests on two connected pillars:

Intelligence

That brings people, data, workflows
and agents together.

Trust

that applies security, identity, governance
and management consistently across them.

For many organisations, that foundation already exists in Microsoft 365 – Copilot brings AI into the flow of work, while the broader platform provides the controls to scale responsibly. And because AI itself keeps changing, that foundation must stay model-flexible – able to bring new and evolving models into the same governed platform, so you can match the best model to each scenario without re-plumbing security and governance each time.

In this guide

01

Why scaling AI is now a control problem, not just a productivity one

02

The technical and operational blockers that stall pilots

03

A five-part model for scaling AI securely – the spine of this guide

04

How model choice keeps you flexible without multiplying risk

05

A scored self-assessment to gauge your agent readiness

06

Where to start, and how Microsoft can help

Why AI pilots stall before they scale

You've likely invested heavily in AI experimentation and seen real results: time saved on routine tasks, faster insights, more efficient workflows. Expanding those wins beyond isolated teams is where progress slows.



Gartner has projected that at least 30% of generative AI projects are abandoned after proof of concept – citing poor data quality, inadequate risk controls, escalating costs and unclear business value.

GARTNER, 2024



Four barriers IT Leaders hit when they try to scale

01

Tool sprawl creates complexity

As teams adopt different tools, you accumulate overlapping capabilities, duplicated spend, fragmented experiences and inconsistent controls. What began as experimentation becomes an operational burden that makes cost, security and value harder to manage.

02

Value gets harder to measure

Early pilots focus on individual productivity, but you eventually need to show broader business impact – better processes, faster decisions, consistent adoption and a clear view of which investments actually create value.

03

Governance lags adoption

Business teams want new capabilities immediately; security and compliance leaders need confidence that controls are in place. Balancing the two has become the defining challenge of enterprise AI.

04


The agentic shift raises the bar

As agents begin to perform multi-step tasks, they introduce new demands around identity, access, permissions, monitoring and lifecycle management. You need confidence that today's platform decisions will support tomorrow's agents without opening new gaps in visibility or control.


A new model for scaling AI

The five foundations

The first wave of AI was about experimentation. The next is about running it at scale – keeping visibility, control and flexibility while getting more value from the investments you already own. Five foundations consistently separate the organisations that move from pilots to enterprise value. They're the spine of everything that follows in this guide.



AT A GLANCE



Maximise AI value at scale

With the model in place, the next question is where AI creates measurable value. It's a common misconception that value comes from the technology itself. It comes from how well AI is integrated into the work, data, security and management practices you already rely on – AI is most powerful when it extends your existing systems rather than operating in isolation.

Organisations often start with individual productivity – summarising, drafting, researching, creating. Those are valuable entry points. The larger opportunity emerges when AI transforms workflows and business processes.

Three questions help find where AI and agents can create the greatest impact:

1. Where is knowledge fragmented across systems?

Important information is often spread across applications, documents and teams, creating delays and inconsistency in execution.

2. Where do workflows break across people and systems?

Agents can coordinate tasks, trigger actions and connect systems so processes move forward with less friction.

3. Where is transformation limited to visible work, not the underlying process?

Many organisations chase surface-level productivity gains, while the greatest impact comes from rethinking the back-end processes that shape outcomes.



What it looks like in practice

In Forrester's study of Microsoft 365 Copilot, interviewees described post-meeting recaps completed in about five minutes instead of 30, and a professional-services firm that produced 10% more sales proposals per month while closing deals in 20 days instead of 30.

FORRESTER TOTAL ECONOMIC IMPACT™ STUDY
COMMISSIONED BY MICROSOFT.



That value can start with ready-to-use agents in Copilot – built by Microsoft and trusted partners – then expand through custom agents your own teams create, grounded in your knowledge sources, data, systems and processes. The goal isn't automation for its own sake. It's helping people spend more time applying expertise, making decisions and creating value, inside the Microsoft 365 apps they already use.



Optimise cost and simplify growth

As adoption expands, you face a new operating question: **where is AI actually being used, which investments create value, how are costs evolving and where do capabilities overlap?** Without that visibility, AI growth outpaces your ability to manage spend and connect investment to outcomes.

A platform approach simplifies these decisions by connecting AI adoption to the productivity, security, management and governance investments you already own. Consolidation reduces operational complexity and gives employees a more consistent, familiar experience. It also lets you manage AI spend with real visibility – connecting investment to measurable outcomes, seeing which capabilities drive value and deciding where to scale automation, where human expertise remains essential and how to balance the two.

THE CONVERSATION SHIFTS FROM

*How many AI tools
do we have?*

↓ to

*How effectively are
our AI investments
supporting business goals?*

Innovate with confidence on a future-ready platform

Innovation and control are often cast as competing priorities. In practice, you need both – freedom for employees to explore, and confidence that innovation can scale responsibly. That confidence is the trust pillar in action.



Here's what it looks like in practice

Identity and access

As AI connects more deeply to business data and workflows, extending the identity and access controls you already use is essential – so every experience is grounded in the right permissions, users and context.

Agent governance

As agents take on more, you need visibility into how they operate, what actions they take and how those actions align to policy – enabling innovation with accountability.

Data protection


AI outcomes depend on trusted data. Built-in protection helps maintain confidence in how information is accessed, used and shared across AI experiences.

A future-ready platform

As AI advances, the platform evolves with it – new models, multi-model flexibility, continuously delivered innovation and new agent experiences such as Cowork and Scout. That adaptability lets you adopt what's next without rebuilding your foundation or disrupting work.

Context and intelligence

AI is only as valuable as the context it understands. With Work IQ, Copilot draws on your organisation's data, content and ways of working – so prompts are grounded in relevant context and can produce more useful, personalised and actionable results.



Assess your agent readiness

Use this self-assessment to turn the guide’s principles into a practical readiness conversation across your business, IT, security and transformation leaders.
Score each statement from 1 to 3, then total your score and read your band below.

	1	2	3
Readiness statement	Not started	In progress	Established
Strategy and leadership			
Executive sponsorship and a clear AI vision are established.			
AI and agent initiatives are aligned to measurable business outcomes.			
Business process readiness			
High-value workflows and use cases for AI and agents have been identified.			
Key business processes, dependencies and success measures are understood.			
Technology and data readiness			
Employees can access trusted information when they need it.			
High-value knowledge sources and business data have been identified.			
Security and governance readiness			
Clear controls for data access, protection and compliance are in place.			
Governance policies can be applied consistently across AI and agent experiences.			
Processes exist to manage agent lifecycle, monitoring and oversight.			
Organisational readiness and adoption			
Employees are prepared to use AI effectively, with a plan for ongoing adoption and change management.			
Stakeholders understand how AI and agents will support the way work gets done.			

Score your readiness

Add up your scores across all 11 statements (maximum 33), then find your band:

Your score	Range	Where to focus next
Ready to scale	27–33	Foundations are strong. Prioritise your first agent-enabled scenarios and lock in value metrics – start at Step 4.
Emerging	18–26	The base is forming. Close gaps in governance, data access and success measures before you scale – start at Step 2.
Early	11–17	Begin with strategy, executive sponsorship and identifying high-value workflows – start at Step 1.

Where to start

After assessing readiness, the next step is turning understanding into a plan. Scaling AI rarely begins with buying more tools. It begins with finding where AI can create the greatest impact and where you're ready to scale responsibly.

A practical first sequence:

01

Identify high-value workflows.

02

Assess current AI usage across the organisation.

03

Evaluate governance and security readiness using the platform you already run.

04

Prioritise the first agent-enabled scenarios.

05

Define adoption and value metrics.

06

Build a roadmap for long-term scale.

Organisations that take a deliberate, platform-driven approach are better positioned to balance innovation, governance and business value.

How Microsoft helps

The choices in this guide point to one conclusion: the next generation of work needs more than standalone AI tools. It needs a **connected platform** that brings **productivity, business data, security, identity, governance, management and AI together** in one foundation.

Microsoft 365 E7 is that foundation in a single integrated solution. It combines four building blocks:

- Microsoft 365 E5 for productivity, security and management;
- Microsoft 365 Copilot for AI in the flow of work;
- Microsoft Agent 365 for agent management; and
- Microsoft Entra Suite for advanced identity and access.

In short: E7 is E5 plus the AI, agent and identity layers you need to scale – governed as one.





The returns can be real when the foundation is right.



In a Forrester Total Economic Impact™ study commissioned by Microsoft, a 25,000-person composite organisation modelled a 116% three-year return on its Microsoft 365 Copilot investment, with payback in under 11 months and an average of roughly nine hours saved per user each month.

FORRESTER, 2025

Notably, that composite first invested in data-hygiene and permissioning work before rollout – a reminder that AI value depends on governed, well-permissioned data, which is exactly what a platform approach is designed to provide.

As model choice expands, Microsoft's advantage isn't simply access to more models. It's the ability to adopt model diversity within the same trusted platform you use to govern people, data, apps and agents – without repeated effort across evaluation, tuning and maintenance.

Rather than stitching together separate AI, identity, security, governance and agent-management tools, **E7 connects them, keeping people in control as AI moves work forward.** For teams that want to go further, Copilot Studio lets them build and manage agents tailored to their own processes, knowledge and systems.

With Microsoft 365, intelligence and trust come together in the flow of work. You can evaluate and use the right models, govern securely and scale AI with confidence as models and agents keep evolving – without a separate operating model for each one. That's what it means to put AI in the flow of human ambition: keeping people and their ambition, at the centre of how work moves forward.



Take the next step

Netskope, Cloud and Threat Report: Shadow AI and Agentic AI, 2025 – majority of enterprise AI use occurs through personal, unmanaged accounts

Gartner, 'Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept by End of 2025', July 2024.

Forrester Consulting, The Total Economic Impact™ of Microsoft 365 Copilot, commissioned by Microsoft, 2025 (116% three-year ROI; payback under 11 months; ~9 hours saved per user per month); and New Technology: The Projected Total Economic Impact™ of Microsoft Copilot for Microsoft 365, commissioned by Microsoft, April 2024 (meeting-recap and sales-proposal examples).

Note: external (non-Microsoft) figures should be confirmed against CELA's preferred citation list before publication. A named ME7 customer story may be substituted for the practice example.



Try Microsoft 365 Copilot today

See how AI in the flow of work helps your teams create, collaborate and move faster with more context and control – start a Copilot chat and put it to work on a real task today.



Scale AI with confidence with Microsoft 365 E7

Bring together productivity, security, identity, AI and agent management in one integrated solution designed to help you scale AI with greater confidence, control and business value.

