

AI Readiness: A Practical Guide to Adopting AI Securely and Measurably

How small and mid-sized organizations can govern shadow AI, prepare Microsoft 365 for Copilot, strengthen cyber resilience, and prove business value

By AdvancedBITS | Business IT Solutions • Updated October 2026 • About 12 minutes to read

Quick answer: what is AI readiness?

AI readiness is having governed AI use, organized and correctly permissioned data, strong identity security and recovery, trained people, and defined success measures in place before a business adopts AI tools such as Microsoft 365 Copilot. Employees are already using AI, attackers are too, and many organizations cannot yet show it improves the business. The fix is the same each time: govern the tools, clean up the data and permissions behind them, secure identity and recovery, train people, and measure results against a baseline.

Four pressures look like separate problems: unapproved AI tools, Microsoft 365 Copilot readiness, AI-enabled cyber risk, and proving ROI. In practice they rest on one foundation. This guide brings them together into one readiness model, one roadmap, and one set of questions for leadership.

What Does the Research Say About AI Risk and ROI?

Independent research shows employees adopting AI faster than organizations govern it, the cost of weak security basics rising, and individual productivity gains that are not yet turning into business results.

Finding	Source	Why It Matters
78% of AI users bring their own AI tools to work, and 60% of leaders say their company lacks a vision and plan for AI.	Microsoft and LinkedIn, 2024 Work Trend Index	AI activity exists before AI strategy does.
69% of surveyed cybersecurity leaders suspect or have evidence of prohibited public AI use. Gartner predicts more than 40% of enterprises will have shadow AI security or compliance incidents by 2030.	Gartner, November 2025	Security leaders already see it.
39.7% of data movements into AI tools involve sensitive data, and 32.3% of ChatGPT use is through personal accounts.	Cyberhaven Labs 2026 press release; report page (222 companies)	Sensitive data entry is routine.
One in five organizations reported a breach tied to shadow AI, costing \$670,000 more on average than breaches with little or no shadow AI. 63% of breached organizations had no AI governance policy.	IBM Cost of a Data Breach 2025, via Cybersecurity Dive	Unmanaged AI has a measurable price.
One in four malicious breaches was AI-enabled, averaging \$6 million versus a \$4.99 million global average. AI and automation in security operations cut breach costs by almost \$2 million.	IBM Cost of a Data Breach 2026	AI raises the stakes for attackers and defenders.
Reported cybercrime losses reached nearly \$21 billion in 2025. The FBI's first AI section counted 22,364 complaints and nearly \$893 million in losses.	FBI 2025 Internet Crime Report	AI-enabled fraud is now its own category.

Finding	Source	Why It Matters
69% of small and mid-sized businesses refused ransom demands because they had reliable backups.	Cyber Readiness Institute, summarizing Verizon 2026 DBIR	Tested backups change ransomware outcomes.
Eight in ten say AI improved their own productivity, but only 37% attribute any EBIT impact to AI, unchanged from the prior year.	McKinsey, State of AI 2026	Individual gains do not become business results on their own.
A preliminary MIT NANDA report found 95% of organizations in its study saw no measurable return. Wharton and GBK's 2026 survey of 1,416 U.S. enterprise respondents found 80% of enterprise leaders report positive AI returns, but only 28% say those returns are strongly positive.	MIT NANDA, 2025 ; Wharton-GBK, Oct. 2026	Results are mixed, so measure your own.

What Is the AdvancedBITS AI Readiness Model?

AI readiness is not one checkbox. It is six connected areas; each covered in the sections that follow.

Area	Question to Ask	Risk if Ignored	Better Approach
Governance and policy	Which AI tools, data types, and uses are approved?	Shadow AI. Sensitive data leaves approved systems.	Acceptable-use policy, approved-tool list, and a named owner.
Data and permissions	Is information organized, current, and permissioned correctly?	AI surfaces overshared or outdated content.	Clean up Microsoft 365, SharePoint, Teams, and OneDrive access before broad rollout.
Identity and security	Are users, devices, and accounts protected?	A compromised account reaches everything AI can reach.	MFA, Conditional Access, endpoint protection, and monitoring.
Resilience and recovery	Can data and operations be restored, and are payment requests verified?	AI-enabled attacks and failures cause lasting disruption.	Tested backup, incident plans that include AI, and out-of-band verification.
People and review	Do employees know the rules and where humans must review?	Overtrust, misuse, and inconsistent adoption.	Role-based training and human review of high-impact output.
Measurement	How will leadership know AI is working?	ROI becomes a feeling instead of a result.	Baseline, scorecard, and stop-or-scale decisions.

The model draws on the four functions of version 1.0 of the voluntary [NIST AI Risk Management Framework](#) (Govern, Map, Measure, and Manage) and applies them in a form that small and mid-sized businesses can act on. The most common mistake is buying licenses first and hoping value follows. Run a readiness review, then start with one workflow and one team.

What Is Shadow AI, and Why Is Banning It Not Enough?

Shadow AI is the use of AI tools, accounts, browser extensions, assistants, or agents outside the organization's approved governance process. Typical examples are public chatbots on personal logins, AI note takers that join meetings, and AI features adopted inside applications without review. It usually comes from motivated employees rather than bad actors, which makes it both a risk signal and a demand signal.

A ban rarely removes the demand, and it makes the activity harder to see. [Gartner advises CIOs](#) to define clear policies for AI tool use, audit regularly for shadow AI, and include generative AI in SaaS risk reviews. AdvancedBITS recommends controlled enablement: provide approved tools that meet real needs, define what data may never enter unapproved tools, train by role, review vendors, and give employees a simple way to request new tools. [Microsoft's shadow AI guidance](#) likewise puts discovery and monitoring ahead of any blocking.

Data Tier	Examples	Rule of Thumb
Low risk	Published web content, public marketing copy, generic research questions with no company or client details.	Allowed in approved tools. Verify accuracy.
Restricted	Internal procedures, non-sensitive drafts, internal meeting notes, anonymized data.	Approved enterprise accounts only, with review rules.
Prohibited in unapproved tools	Client and employee records, non-public financials, credentials, source code, contracts, regulated data.	Never enter into personal or unapproved AI tools.

Avoid vague policy such as “use AI responsibly.” Give practical examples of allowed, restricted, and prohibited use, like the tiers above, and train by role. Polished output is not correct output, so require human review before client, financial, legal, or HR decisions.

How Do You Prepare Microsoft 365 for Copilot?

Microsoft 365 Copilot readiness is the work of preparing a tenant's permissions, sensitive data, identity controls, and users before Copilot is turned on. Copilot answers using content each user can already access. That is sound security design, but it means a file shared too broadly years ago becomes easy to find with a plain-English question. Copilot does not create oversharing. It reveals it. [Microsoft's deployment guidance](#) has three steps: remediate oversharing, set up guardrails, and meet regulations. Common oversharing signals include:

- Sites or files open to “Everyone except external users” or company-wide groups.
- Ownerless or inactive sites and abandoned Teams.
- Broken permission inheritance and long-lived external sharing links.
- Public site or team privacy and broad default sharing links.
- Sensitive files that were never labeled.

Microsoft provides built-in tools to find these: the [SharePoint Advanced Management Content Management Assessment](#) (Microsoft suggests rerunning it every 30 days), [data access governance reports](#), and Purview data risk assessments, with [DLP and sensitivity labels](#) as guardrails. SharePoint Advanced Management is included with Copilot licenses. Microsoft also recommends interim protections to reduce exposure before cleanup is finished, so a focused pilot does not have to wait for perfection. Do not skip the cleanup because sites look manageable today. Prioritize high-risk sites, broad permissions, and sensitive libraries first.

How Does AI Change Cyber Risk and Business Resilience?

AI does not replace the security fundamentals. It raises the cost of neglecting them. Phishing is more convincing, voice and video can be faked, and AI tools and agents become new places for data and access to leak. The most striking example is Arup: in January 2024 a finance employee joined a video call in which [every other participant was an AI-generated deepfake](#) and authorized 15 transfers totaling about US\$25.6 million. Arup's CIO called it [technology-enhanced social engineering](#) rather than a system compromise, which is why procedures are part of resilience. Useful controls, offered as AdvancedBITS guidance informed by FBI advice, include:

Control	How It Works
Out-of-band verification	Confirm payment, bank-detail, and credential requests through a second channel, using a contact already on file, never one supplied in the request.

Control	How It Works
Dual approval	Require two approvers for wire transfers, new payees, and banking changes above a set threshold.
No approvals by voice or video alone	A call, video meeting, or voicemail is never sufficient authorization to move money.
Pause rule	Give staff explicit permission to stop and verify urgent or confidential requests.
Report fast	Define who staff call first. For fraud involving funds, contact the bank immediately and file at ic3.gov.

- **Identity:** [Microsoft Research](#) found MFA reduced account-compromise risk by 99.22% across the population studied. It is necessary but not sufficient, so pair it with Conditional Access and monitoring.
- **Patching:** In Verizon's 2026 data, [vulnerability exploitation was the top way in](#) (31% of breaches), with a median of 43 days to fully patch.
- **Backup:** Backups only help if restores are tested against real business processes and Microsoft 365 data is covered, not just servers.
- **Governance:** Do not treat AI as separate from security. Include AI tools, agents, and vendors in security reviews and monitoring.
- **Incident response:** Add AI scenarios to your plan, such as sensitive data entered into an unapproved tool, a deepfake payment request, or an AI agent acting unexpectedly.

How Do You Measure AI ROI and Prove Business Value?

Usage is not value. As the research table shows, individual productivity gains are not turning into financial results for most organizations. McKinsey's [earlier analysis](#) found that redesigning workflows was the attribute most strongly linked to EBIT impact from gen AI among 25 tested. The answer is to measure outcomes in one specific workflow: time saved, quality, service speed, risk reduced, capacity, and revenue supported. A 90-day pilot looks like this:

- **Define** one recurring workflow with clear pain and ownership.
- **Baseline** current time, quality, volume, backlog, or response before the pilot, and track outcomes rather than logins. The UK DWP's own [Copilot evaluation](#) lists the absence of a pre-trial baseline as a limitation.
- **Prepare** permissions, data sensitivity, backup, identity, and approved tools.
- **Train** the pilot group by role, with examples and review rules.
- **Measure** results against the baseline using metrics agreed in advance.
- **Decide** whether to scale, revise, or stop.

Count time saved conservatively. In the UK government's [20,000-person Copilot trial](#), participants reported saving 26 minutes a day, but the figures were self-reported. As a hypothetical, if ten users each save two hours a week and only half is redirected to higher-value work, count 10 hours, not 20. Ninety days is deliberate: Microsoft's [WorkLab research](#) notes it can take up to 11 weeks to build the AI habit.

What Does a Six-Month AI Readiness Roadmap Look Like?

Phase	Focus	Outcome
Months 1 to 2: Discover and decide	Inventory AI tools in use, review identity controls and Microsoft 365 permissions, identify sensitive data, assess backup coverage, set interim data rules, choose one low-risk pilot workflow, and capture its baseline.	A realistic view of exposure and a measurable pilot target.

Phase	Focus	Outcome
Months 3 to 4: Protect and pilot	Publish the acceptable-use policy, approve initial tools, fix the highest-risk oversharing, strengthen MFA and Conditional Access, train a pilot group (including AI-enabled scams), validate restores, and run the pilot with human review.	Controlled adoption with evidence.
Months 5 to 6: Govern and scale	Review vendor risk, add AI scenarios to incident response, compare pilot results to the baseline, expand only if value is proven, and set quarterly governance and scorecard reviews.	A repeatable model for responsible AI growth.

What Should Leaders Ask Before Approving AI Use?

- Which AI tools are employees using today, and do we offer an approved option that meets real needs?
- What sensitive data could be entering them, and which data is prohibited?
- Are Microsoft 365 permissions ready for AI-assisted search?
- Do MFA and Conditional Access apply consistently, and can we restore data after deletion or ransomware?
- Does our incident plan cover AI data exposure and deepfake payment fraud?
- Which workflow are we improving first, and what is its baseline?
- Who owns AI governance and reviews outcomes every quarter?

How Does AdvancedBITS Help With AI Readiness?

AdvancedBITS helps organizations build the secure, governed, and measurable foundation that responsible AI adoption requires. Our role is not simply to promote AI tools.

Area	AdvancedBITS Support
AI readiness and governance	Readiness assessments, shadow AI discovery, acceptable-use policy support, and tool and vendor review.
Microsoft 365 governance	Tenant readiness review, SharePoint, Teams, and OneDrive permission cleanup, retention, and document-structure guidance.
Identity and cybersecurity	MFA, Conditional Access, managed security, endpoint protection, email security, and monitoring.
Backup and recovery	Microsoft 365 backup, restore planning, business continuity, and recovery validation.
Incident readiness	Incident response planning, escalation guidance, and tabletop exercises that include AI scenarios.
Training and measurement	Role-based training, 90-day pilot design, baseline and scorecard development, and adoption review.

Frequently Asked Questions About AI Readiness

What is AI readiness?

AI readiness is having governed AI use, organized and correctly permissioned data, strong identity security and recovery, trained people, and defined success measures in place before adopting AI tools such as Microsoft 365 Copilot. It is a planning and governance question, not a software purchase.

What is shadow AI?

Shadow AI is the use of AI tools, accounts, extensions, or agents, often public chatbots on personal accounts, without IT, security, or management approval. Leadership cannot see what data is entered, where it goes, or how it is retained, which creates security and compliance risk.

Why prepare before turning on Microsoft 365 Copilot?

Copilot can surface any content a user can already access, so files shared too broadly become easy to find with a simple question. Reviewing SharePoint, Teams, and OneDrive permissions first keeps AI-assisted search from showing information to the wrong people.

What should an AI acceptable-use policy include?

It should name approved, restricted, and prohibited tools, list the data types that may never be entered into unapproved AI, state when human review is required, explain how employees request a new tool, set vendor review and training expectations, define incident reporting, and name an owner for AI governance.

How do you measure the ROI of AI?

Pick one workflow, capture a baseline before the pilot, run a trained group for about 90 days, and compare results against metrics agreed in advance. Track outcomes such as time, quality, service speed, risk, and capacity, not just usage, and count only time that is actually redirected to higher-value work.

Do most AI pilots really fail?

No single failure rate applies to every business. A preliminary 2025 MIT NANDA report found most pilots in its study produced no measurable return, while Wharton and GBK's 2026 survey found most enterprise leaders report positive returns. Results depend on workflow fit, so measure your own pilot against a baseline.

Do Not Let AI Become an Unmanaged Experiment

AI should not be ignored or adopted carelessly. Govern the tools, prepare the data, secure identity and recovery, train the team, and measure the value before expanding. AdvancedBITS helps clients take that disciplined path.

AdvancedBITS | Business IT Solutions

Secure. Compliant. Outcome-Focused.

sales@advancedbits.com | www.advancedbits.com | 561-300-3015

Sources and Further Reading

Statistics and guidance come from the sources below. Read figures as directional benchmarks: Verizon data is cited via third-party summaries, the MIT NANDA report is preliminary, and frameworks and examples are illustrative.

- [Microsoft and LinkedIn 2024 Work Trend Index](#); [Gartner GenAI blind spots \(Nov. 2025\)](#); [Cyberhaven Labs 2026 report](#)
- [IBM Cost of a Data Breach 2025 \(via Cybersecurity Dive\)](#) and [2026 newsroom summary](#)
- [FBI 2025 Internet Crime Report \(April 6, 2026\)](#) and Verizon DBIR 2026 via [SecurityWeek](#), [Cyber Readiness Institute](#)
- [Microsoft Research MFA study \(May 2023\)](#); [NIST AI RME](#); Arup case: [Insurance Business](#), [Adaptive Security](#)
- [McKinsey: The State of AI in 2026 \(August 25, 2026\)](#) and [The State of AI, March 2025 \(PDF\)](#)
- [MIT Project NANDA: The GenAI Divide \(July 2025, preliminary, PDF\)](#) and [Wharton-GBK 2026 AI Adoption study \(Oct. 13, 2026\)](#)
- UK Copilot evidence: [GDS trial \(2025\)](#), [DWP evaluation \(2026\)](#), [Microsoft WorkLab](#)
- Microsoft Learn: [Copilot foundation](#), [SharePoint Advanced Mgmt](#), [access reports](#), [Purview DLP](#), [shadow AI](#)