

AI at Work: 10 Settings That Protect Your Business

Your team is already using ChatGPT, Copilot, and similar tools. The question is whether that use is protecting your business or quietly exposing it. These ten items are where to start.

SETTING / FEATURE	WHY IT MATTERS	RISK IF NOT IMPLEMENTED
1. Use approved AI tools on business accounts Only company-approved tools (ChatGPT Team/Enterprise, Microsoft 365 Copilot) with accounts the company owns and manages.	Business plans come with data-protection terms, admin controls, and the ability to remove access when someone leaves. Personal and free accounts have none of that.	Employees paste client data into tools you cannot see or control. When they leave, the account and its history leave with them.
2. Turn off model training on your data Disable any setting that lets the vendor use your conversations to improve its models. Business tiers exclude training by default; confirm it.	Anything typed into a tool that trains on your data can be retained and reused. There is no way to pull it back out.	Contracts, pricing, client names, and internal notes become part of a model anyone can query.
3. Require MFA and single sign-on on every AI account AI logins tied to your Microsoft or Google identity, with multi-factor authentication enforced.	AI chat history is a complete record of how your business thinks: deals, problems, people. A stolen password exposes all of it at once.	One phished login gives an attacker months of conversations. Cyber insurance claims are routinely denied when MFA was not in place.
4. Define what never goes into a prompt A short, written list: passwords, Social Security numbers, card numbers, patient or client records, payroll, unreleased financials.	Once data is sent to an AI tool it has left your control. Many of those data types carry legal or contractual obligations (HIPAA, PCI, client NDAs).	A compliance violation and breach notification triggered by a well-meaning employee trying to save ten minutes.
5. Put a plain-language AI use policy in writing One page. What is allowed, what is not, who to ask. Every employee reads and signs it.	Your team is already using AI. A policy sets expectations, protects the company, and gives you something to point to in an audit or dispute.	No accountability, no consistent behavior, and no defensible position if something goes wrong.
6. Review AI output before it leaves the building A human checks anything AI helped write before it goes to a client, into a contract, or into a financial decision.	AI tools produce confident, polished answers that are sometimes wrong. The tone does not change when the facts are off.	Incorrect advice sent under your name, a bad number in a proposal, or a legal citation that does not exist.
7. Control what your AI tools are allowed to reach Decide which systems an AI tool may connect to (email, files, CRM, accounting) and clean up who has access to what before turning it on.	An AI assistant can read everything the person using it can read. If file sharing is loose, the AI makes that problem visible to everyone.	An employee asks a simple question and the AI hands back HR files, payroll, or the owner's private documents.
8. Lock down company-wide settings in the admin console Set the rules once in the administrator panel: data retention, which features and add-ons are allowed, who can build custom assistants, and company instructions every user inherits.	Settings made at the company level apply to every employee and cannot be switched off by an individual. Settings left to each person will drift.	Fifteen employees running fifteen different configurations, with the safest ones depending on whoever happened to read the manual.
9. Verify the request, not the voice Any request to send money, change bank details, or share a password gets confirmed by calling the person back on a number you already have. No exceptions, even for the owner.	AI can now clone a voice from a few seconds of audio and write an email that sounds exactly like your boss. The old warning signs (bad grammar, odd phrasing) are gone.	A wire transfer sent to a criminal because the call sounded like the CEO, with no way to get the money back.
10. Back it all with 24/7 monitoring and response A security operations center (SOC) watching your accounts, devices, and logins around the clock, with people who act on alerts, not just a tool that sends them.	AI has made attacks faster and cheaper to run, at machine speed, at 2 a.m. on a Saturday. The settings above lower your risk; someone still has to catch what gets through.	An attacker is inside for days or weeks before anyone notices. By then the data is gone and the recovery bill is many times what monitoring would have cost.

Not sure which of these are in place today? Find out in minutes at vitalsecured.com