Guidelines for Working Safely with Artificial Intelligence in the Workplace

PJ Networks
Computer Services

Prepared by PJ Networks Computer Services

Web: www.pj-networks.com • Email: info@pj-networks.com

AI Security Guidelines • 2026 Edition

Introduction

As more and more workplace tools and solutions are being integrated with Artificial Intelligence features and capabilities, it is easy for people to believe that these solutions are more or less equally safe and effective – but that is far from the truth. Not only do they vary widely in their accuracy and capabilities, but some of them have glaring weaknesses in their safety protocols that literally make them dangerous to work with unless you fully understand how that particular solution works.

It would be impossible to write up a single set of AI usage guidelines that covers every single security risk and warning needed to keep an end user (you) safe, since the capabilities of AI are evolving on a daily basis and more solutions are available every time you look. Instead, what we are offering here are general guidelines for safely using Artificial Intelligence in your business environment, although most of this advice also applies to the personal use of AI.

If you take the time to read through these guidelines, we believe that they will make sense to you, and therefore they will be easy to remember. When you were a child you were probably told to never stick a knife or a fork into an electrical outlet – and why. Once you understood the "why" part, you did not have to be told again every time you were near an outlet. You understood the danger.

We hope that these guidelines will clearly explain the dangers of using AI in the workplace in a way that will make sense and stay with you as you begin to explore and work with AI solutions that will hopefully benefit both your business and your personal life.

It doesn't have to be scary or confusing – most of it is just common sense.

RULE #1: You should be aware of every single AI model, tool, or agent that is being used within your work environment, paid or otherwise.

The Corner Cubicle Risk

In many offices, employees seeking to increase their productivity or score favor with management are privately using free, public AI platforms in their 'corner cubicles' without organizational approval or knowledge. For instance, an employee might upload sensitive, unredacted corporate records—such as draft profit-and-loss reports, strategic business plans, or customer contact sheets—into a public chatbot to generate an analytical presentation. In doing so, they are unknowingly publishing confidential company assets to external servers, exposing the company to data leaks, intellectual property theft, and regulatory non-compliance.

Important Security Notice: Treat public AI engines as public bulletin boards. Unless a validated enterprise contract and safety protocol are explicitly in place, assume that anything entered into an AI session becomes part of the public domain and is instantly accessible to the model's parent corporation.

The Golden Rule of AI Privacy: Never upload or share anything with an AI model that you would not feel completely comfortable posting publicly on Facebook or LinkedIn. If you would not publish your Social Security Number, bank routing numbers, proprietary designs, or financial records on a public forum, then don't put them into an unvetted AI session.

Understanding Models: Public vs. Licensed AI

To integrate AI safely, organizations must understand the fundamental structural difference between public (free) models and licensed (enterprise) models. The primary rule of thumb in modern computing remains true: nothing in life is free. When using a free version of an AI chatbot, the user agrees to terms of service that grant the developer permission to retain conversation history and actively train their models on all inputs.

The Free Model Trap Consider a real-world warning scenario: a local business in Charlottesville, Virginia, uses a free public AI tier to draft a comprehensive, proprietary business model and expansion plan to secure a bank loan. By inputting their manufacturing details and financial projections into the free tool, the business effectively grants the AI platform permission to utilize that proprietary structure.

Example:

To understand the real-world consequences of this agreement, consider a proprietary business plan drafted by a widget manufacturer in Charlottesville, Virginia. If the manufacturer spends ten hours using a free AI system to refine their confidential financial projections and market strategy, that data becomes part of the AI's public training set. If a competitor later logs in and asks the same AI tool to generate a widget business plan for the Charlottesville area, the AI can literally drop the first manufacturer's exact, proprietary plan onto the competitor's screen. The AI is not magic; it has simply regurgitated the proprietary details the first user handed it.

Licensed Enterprise Protection

Licensed enterprise tiers (ranging from standard user licenses up to advanced professional accounts like Google Gemini Advanced/Ultra) operate under strict commercial agreements. Under a paid enterprise license, AI developers guarantee that conversation history is fully encrypted, data is isolated, and inputs are strictly blocked from training public models. These paid contracts protect your intellectual property, making them safe for internal operations drafts, approved budget analysis, and strategic modeling.

Feature / Capability	Free Public AI Tiers	Licensed Enterprise AI Tiers
Data Privacy & Encryption	Minimal or none. Conversations are stored on public servers.	Full enterprise-grade encryption and data isolation.
Model Training Use	Yes. Platform actively trains on all inputs and outputs.	Strictly blocked. Inputs are never used to train public models.
Security Guarantees	None. No compliance or contract-based data protections.	Backed by enterprise/university commercial safety contracts.
Best Use Cases	General, public research (e.g., planning a vacation, editing public text).	Internal operations, financial drafting, strategic plans.

Protecting Your Data (as much as possible)

When using major AI platforms, the baseline rule for data privacy comes down to the specific type of account you hold. Free tiers and individual paid subscriptions (such as ChatGPT Plus, Gemini Advanced, or Claude Pro) almost always default to using your conversations to train their future models. Conversely, enterprise, team, and commercial workspace licenses are specifically designed for business compliance and generally exclude your data from model training by default. If you are using any consumer-level account—whether free or paid—you must proactively lock down your settings.

ChatGPT (Free & Plus Tiers) Both the free version and the individual $20/month Plus subscription automatically opt you into data training and retain your chat history indefinitely. To stop this, open your settings, select **Data Controls**, and toggle off "Improve the model for everyone." For highly sensitive, ad-hoc queries, you can also use Temporary Chats; these conversations are excluded from training and are permanently deleted within 30 days. (Note: ChatGPT Enterprise and Team accounts are opted out by default).

Google Gemini (Free & Advanced Tiers) By default, your interactions with Gemini are saved, used for training, and may be subject to human review to improve the service. To opt out, navigate to the **Gemini Apps Activity** page via the settings menu and turn activity tracking off completely. If you use Gemini Live for voice or video features, look for the separate audio and screenshare settings within the Activity menu and disable those as well. (Note: Google Workspace commercial accounts with Gemini enabled generally protect your data by default).

Anthropic Claude (Free & Pro Tiers) While Claude is often considered more privacy-focused and automatically purges data after 30 days on consumer tiers, both the free and individual Pro versions still default to allowing data sharing for model training. Open your profile settings, go to **Privacy or Model Training**, and ensure the toggle to improve AI models is switched off. (Note: Claude Team and Enterprise accounts exclude your data from training by default).

Microsoft Copilot (Consumer & Pro Tiers) Unless you are logged in with a secure corporate Microsoft 365 commercial license, Copilot will use your text and voice inputs for training. Open the **Privacy** section under your profile icon and toggle off model training for both text and voice interactions.

Regardless of the platform or the price tag, no cloud-based AI chatbot is completely private or end-to-end encrypted. Data that has already been absorbed into an algorithm's training set cannot be cleanly extracted or deleted. Taking a few minutes to verify your account tier and actively disable data sharing is the single most effective way to ensure your proprietary information remains entirely yours.

RULE #2: Always take a moment to check your data sharing and privacy settings before you start using any AI tool or resource.

Integrated AI Solutions Like Microsoft 365 Copilot and Google Workspace Gemini

This is where the architecture fundamentally shifts. When an organization deploys an integrated AI assistant directly into their tech stack (such as Microsoft 365 Copilot or Google Workspace Gemini), **the AI no longer relies solely on what the user types into a chat box.**

- **Inherited Permissions:** The AI model is granted access to every single piece of data the employee has the security permissions to view.
- **The Deep Crawl:** If the employee asks the integrated AI to "summarize the latest financial reports," the AI will instantaneously crawl their Outlook emails, personal OneDrive files, shared Teams chats, and every single SharePoint site they are authorized to access to compile that data.
- **The Hidden Vulnerability:** If the company's IT administrators have not strictly locked down internal permissions—meaning standard employees accidentally have read-access to a sensitive HR folder or an executive SharePoint drive—the AI will find that data and serve it directly to the employee, exposing internal vulnerabilities at lightning speed.

RULE #3: An AI tool being used by an employee will generally have access to everything that the employee has access to – check your company file and folder permissions.

Example: **An employee using Microsoft Copilot might use a query asking for numbers related to a company's expenses. If that employee unknowingly has permission to access the company's payroll information, Copilot will search that information and include it in the report if it is relevant to the task request.**

Critical Risk Zones & Screen Awareness

To protect your organization, certain classes of data must be strictly barred from public AI platforms. Employee security training should explicitly define these risk zones to avoid catastrophic exposures.

Data Categories that Should NEVER be Shared with Public AI platforms :

- **Personally Identifiable Information (PII):** Social Security numbers, dates of birth, home addresses, or patient records.
- **Financial Assets:** Credit card numbers, bank account numbers, routing strings, and active profit-and-loss ledgers.
- **Legal & Vendor Contracts:** Confidential agreements, active strategic plans, proprietary research, or unpublished patents.

The Hidden Threat: Screen-Aware AI Sessions

A newly recognized threat in browser-based AI sessions is that many modern tools are increasingly 'screen-aware'. When an employee launches an active AI session within a web browser, the active web application has the potential technical capability to read other open browser tabs, minimized applications, and active background documents. If an employee is working on sensitive, unencrypted corporate financial files in the background while chatting with a free AI assistant in another tab, those background files are most likely exposed – **yet another reason to avoid free AI tools in the workplace.**

Rule of Thumb: Employees should fully close all sensitive local documents and unrelated browser tabs before starting any AI session. An ounce of prevention is worth a pound of cure.

Mitigating AI Flaws & Cyber Threats

AI tools are not search engines or fact checkers—they are complex, mathematical pattern-guessing calculators. Understanding the inherent flaws and limitations of AI models is essential to protecting business credibility.

The Hallucination Checkpoint

AI systems are engineered to generate responses with high linguistic confidence, which means they can fabricate fictional facts, statistics, and citations that sound entirely plausible. Statistically, 66% of users admit to utilizing AI outputs without verifying them, and over 56% of employees have made professional mistakes due to unverified AI claims. Businesses must enforce secondary verification policies: never accept an AI statement of fact, legal citation, or market statistic without verifying it directly through a trusted, primary source.

The Human-in-the-Loop 80/20 Rule

A vital guideline for secure integration is the Human-in-the-Loop activation curve. AI is incredibly efficient and can get a business draft, code block, or marketing outline to the 80% completion mark almost instantly. However, the final 20% of the effort must be filled by a human being. This final stretch requires rigorous critical thinking, professional judgment, factual validation, and total accountability. If an employee relies on AI to write their output, they remain personally responsible for any errors or biases contained in that output.

Governance: The CAPTURE Framework

To help your employees safely pilot these technologies, PJ Networks recommends implementing the CAPTURE safety framework across all business divisions:

- **C**onfidentiality first: Never enter protected, private, or sensitive company data into public LLMs.
- **A**pproved tools only: Utilize only AI platforms vetted, approved, and licensed by company IT management.
- **P**ilot mode: Keep your hands at the controls. Always act as the pilot, treating the AI strictly as an assistant.
- **T**reat as helper: Understand that AI is a drafting assistant, not a definitive or final source of truth.
- **U**nverified warning: Verify every reference, external link, citation, and factual assertion before distribution.
- **R**eport issues: Promptly report suspected data leaks, anomalous outputs, or suspicious tool behavior to IT.
- **E**ngage security teams: Consult IT and security professionals when AI capabilities and data controls are unclear.

RULE #4: Your business should have an Acceptable Use policy specifically for using Artificial Intelligence, or it should be a part of the main Acceptable Use Policy for your business.

PJ Networks can provide you with a template for either of those policies upon request – free of charge.

The Safest Way To Start Working With AI Models

PAID VERSIONS of Microsoft 365 Copilot and Google Workspace Gemini

By default, Google Gemini and Microsoft Copilot will have access to your company data that is stored in your Google Workspace (for Gemini) or Microsoft 365 OneDrive and Sharepoint (for Copilot). If you are using the paid versions of those tools, then all of the work that you do with them will be limited to your account or anyone else that may have full access to your account. This makes them a relatively safe way to start running queries and building agents that will work with the data that you already have access to without requiring you to open up additional access to outside sources.

EXAMPLE: Using paid Microsoft CoPilot, you can build a prompt and save it as an agent that will look at your Outlook mailbox and Calendar, your Teams messages and Calendar, and your OneDrive folders to determine whether there are any tasks that you were asked to perform or any requests that anybody made of you that have not yet been completed.

The initial prompt might look something like this:

“Using read-only access, look through the last 30 days in my Outlook messages and Calendar, my OneDrive and Sharepoint folders, and my Microsoft Teams messages and identify any requests that have been made of me by others or notes that I have created that would indicate a task that has not yet been completed.”

Please keep in mind that there is a big difference between working with free versions of Microsoft Copilot and Googe Gemini and using the paid versions, so if you want to start safely experimenting with AI tools, we strongly encourage you to sign up for a paid subscription for one of them.

RULE #5: ALL paid AI tools should be owned and licensed by your company, and registered to a company business e-mail address.

Otherwise, when an employee leaves your business, you may lose access to that tool and all of the work that it generated.

Employee Pre-Flight Checklist

1. **Know Your Tools:** Verify that the specific AI platform you are launching is on the company's approved list.
2. **Filter Your Data:** Anonymize all spreadsheets or text inputs by removing bank details, client names, and specific identifiers.
3. **Own the Output:** Review, edit, and sign off on the final work. You are accountable for the quality and accuracy.

Work With a Professional for Advanced AI Needs

Artificial Intelligence can be an amazingly powerful tool and a game changer in achieving a successful outcome as a business. While the simpler tasks can be fairly easy to accomplish on your own with a little practice and a healthy dose of caution, more complicated needs should tackled with the help of an experienced professional.

As a Microsoft partner, PJ Networks can help you to add Microsoft Copilot to your current Microsoft 365 subscription plan, but we do not currently provide AI services or solutions to our clients. However, we can still provide guidance and help to make sure that you understand the risks before you start using Artificial Intelligence in your business.

Pease let us know if you currently are or plan to start using AI in your business so that we can help to make sure that you do that safely. **We recommend engaging a professional consultant** for any AI-based solutions that you want to give access to your critical business data and integrate into your business workflow processes.

For a video overview of these safety guidelines, you can watch our presentation here: AI Cybersecurity Guidelines for Business

AI Training Materials for Beginners

Free training materials to start working safely working with Artificial Intelligence:

Microsoft Copilot: Microsoft Copilot Skilling Center

Google Gemini: Use Google Workspace with Gemini

Introduction to AI: https://course.elementsofai.com/

Securing Your AI Future with PJ Networks

Integrating artificial intelligence into your business operations should never come at the expense of data security. As your managed IT and security partner, PJ Networks Computer Services is prepared to assist your organization with AI Acceptable Use Policies and safe integration practices. Contact our support team at **support@pj-networks.com** to request additional information.