

AI Privacy & Security Checklist

A quick reference for using AI tools safely at work.

☐ Use Approved Tools + Managed Accounts



- Use your organization's approved, managed account.
- A personal paid plan is not automatically approved.

☐ Know What Incognito Mode Does



- Private browsing reduces local browsing traces.
- It doesn't prevent AI providers from retaining your data.

☐ Use Pseudonyms and Minimize Detail



- Check training, chat history, memory, retention, and sharing separately.
- Turning off history does not guarantee deletion.

☐ Confirm Compliance and Approvals



- Consult IT or legal before using AI with regulated data.
- Verify the service and setup meet applicable requirements.

☐ Protect Your Account



- Use strong passwords and multifactor authentication (MFA).
- Sign out on shared devices; avoid confidential work on public devices.

☐ Understand Data Retention + History Settings



- Chat history, training, memory, and retention are separate settings.
- Turning off history doesn't guarantee deletion. Follow company policies.

☐ Be Careful With What You Enter



- Only share sensitive data when the tool and use are approved.
- Never enter passwords, access tokens, or login credentials.

☐ Check Attachments and Connected Access



- Review what files and business systems AI can access.
- Get approval before connecting agents or granting permissions.

☐ Follow Meeting Recording Rules



- Get approval before recording or transcribing meetings.
- Follow notice, consent, and access requirements.

☐ Verify Outputs and Report Issues



- Check AI-generated content before using or sharing it.
- Report accidental disclosures or unauthorized AI activity promptly.

Before you submit:

Check the tool, account, data, and permissions.
If you are unsure, stop and ask.



POWERED BY



AI OPERATIONS & SECURITY

The Secure AI Adoption Playbook

A playbook designed to allow you to capture the full opportunity of AI without leaving your organization exposed to security and compliance risk.

AI is Both the Largest Opportunity and Risk in Your Business Right Now

AI is changing how work gets done. It drafts proposals, analyzes data, summarizes research, answers customers, and writes code, inside the tools people already use every day. Work that once took days now takes minutes, and teams are taking on volume they would have turned down two years ago. That opportunity is real, and the businesses moving first are pulling ahead.

Adoption moved faster than governance. AI entered most organizations through employees, not through procurement or an IT roadmap. Individual teams adopted tools that saved them time, and those tools were quickly connected to email, file storage, the CRM, and other core systems. Because adoption happened outside normal channels, standard controls were never applied. No security review, no scoping of what data the tool could access, and no written policy on acceptable use.

That gap creates both security and compliance exposure. Client records, health information, payment details, and credentials remain regulated regardless of where an employee pastes them. Unsanctioned AI activity puts that data outside the organization's visibility and outside its control, and the exposure typically surfaces late: during an incident, an audit, or a client review.

The answer is not to slow AI adoption. It is to adopt AI responsibly, with the visibility and guardrails your business needs.

Visibility shows which tools are in use, who is using them, and what data is moving. Guardrails set what is acceptable and enforce it at the moment work happens, before data leaves. Nobody can advise on what they cannot see, and nobody can protect what they cannot control.

This playbook is how you capture the full opportunity of AI without leaving your organization exposed. What to look for, what to control, and where to start.

90%

of SMBs are actively using and experimenting with AI.

SOURCE: Q2 2026 PAX8 SMB AI PULSE REPORT

1 in 5

organizations experienced a breach due to a security incident involving shadow AI.

SOURCE: IBM COST OF A DATA BREACH REPORT

\$5.4M

Average cost of a breach involving shadow AI.

SOURCE: IBM COST OF A DATA BREACH REPORT 2026

AI Does Not Just Add New Threats, It Expands the Terrain.

01. Shadow AI

The AI applications and agents your people use without approval, most often through personal accounts. Client records, financials, and internal documents get pasted in with no record of what was shared, no control over how it is stored or whether it trains a model, and no way to pull it back.

66% of employees have used AI at work against company policy

02. Data Leakage

Employees share personal information, client records, and confidential business material with AI applications every day, and none of it is logged or reviewed. That is a compliance violation under HIPAA, PCI, and state privacy law before anything is ever breached. It is also the easiest way for a departing employee to take your data, and if something does go wrong, you have no record of what left.

48% of employees have input confidential data into public AI services.

03. Third-Party & Vendor Risk

Every AI application you run is operated by another company. Your data moves onto their infrastructure, is retained under their terms, and is protected by their controls. When a client or regulator asks where their information went, the answer has to come from you.

12+ AI applications and agents in use per company

04. Rogue AI Agents

AI agents no longer just answer questions, they take action. They send, edit, and reach into other systems through permissions granted long after the application was approved. An agent cleared for drafting text can end up with working access to systems nobody reviewed, acting at machine speed with no one watching.

93% of agent permission prompts are approved by users

Our Approach to Secure AI Adoption

See what is running, secure what matters, and drive adoption that is right for your business.

See.

Complete Visibility

Discover every AI application and agent running across your business, including the ones nobody approved, in a single view.

Secure.

Real-Time Enforcement

Enforce your AI usage policy as it happens, blocking shadow AI, preventing data leakage, and holding autonomous agents inside the limits you set.

Steer.

Guided Adoption

Adopt the tools that genuinely fit your business, with a standing advisor on what to take on next and on what terms.

Ten Questions to Answer in Your First 30 Days

See: What is in Use

- ☐ Which AI tools and agents are in use across the business, including those adopted without approval?
- ☐ Are employees accessing these tools through corporate accounts or personal credentials?
- ☐ What company information has already been entered into AI tools, and by whom?

Secure: Where the Exposure Lies

- ☐ What systems and records can each tool access, and is that access limited to what the role requires?
- ☐ Have any integrations or connections been added since each tool was approved?
- ☐ Do employees know which categories of information may never be shared with an AI tool?
- ☐ If a tool or its vendor were compromised, could access be revoked the same day?

Steer: Who is Accountable

- ☐ Which AI-generated work may be acted upon without human review?
- ☐ What is the process when an employee requests a tool that is not yet approved?
- ☐ Could you demonstrate to a regulator, client, or insurer which AI tools have processed their data?

How We Deliver It.

Just Solutions delivers AI managed services built on KIPIO, an AI operations and security platform purpose-built for MSPs and the businesses they support.

KIPIO enforces your policy at the moment an AI tool is used, identifying which applications are in use, what is being shared with them, and which account was used to sign in, so an unapproved tool reached through a personal login is as visible as a sanctioned one. Each use is evaluated against the data classifications you have defined, and the resulting audit trail belongs to your business rather than to any individual manager.

1 DISCOVER

We assess every user, device, and browser, and return a complete inventory of the AI in use across your business. For each tool we record what data it can reach, who is using it, and what it has been connected to since approval.

2 SET POLICY

We tailor a proven acceptable use policy to your business, defining which AI tools your people may use and which categories of data may never be shared with them. We also agree how exceptions are requested, so the policy still holds the next time a new tool appears.

3 ENFORCE

Your policy is applied at the point of use, so each action is blocked, warned, or allowed as it happens rather than reviewed after the fact. We keep each tool's access limited to what the role requires, and revoke it on demand if a tool or its vendor is compromised.

4 MANAGE

We maintain an audit trail of AI use that will stand up to a regulator or an auditor, and review your inventory and policy each quarter as tools and people change. Just Solutions also advises on which AI tools to adopt next, acting as your chief AI strategy and security officer.

Secure Your Organization's AI Adoption.

justsolutions.com



POWERED BY



AI Acceptable Use and Data Security Policy

For [Organization Name]

Provided by Just Solutions, Inc. for AI webinar attendees

AI is impressive, but it still takes real intelligence to use it well.

Purpose of this template

This editable sample helps an organization permit useful AI work while protecting its information, customers, and people. Once customized and adopted, it establishes which tools, accounts, data, and activities are allowed, who approves exceptions, and how employees report concerns.

Sample only. This document is educational material and is not legal advice. It does not establish compliance with any law, regulation, contract, or security standard. Have qualified counsel and your IT, security, privacy, and HR leads review the customized policy before adoption and when obligations change.

Document control

Field	Customize before adoption
Policy owner	[Name or role and contact]
Approving authority	[Executive or governing body]
Version and effective date	[Version] [Effective date]
Review date	[Next review date] At least annually
Approved tools register	[Internal location or link to Appendix A]
Report an incident	[Security contact and urgent reporting channel]
Related policies	[Acceptable use, classification, privacy, records retention, incident response]

Before publishing your policy

Replace every bracketed placeholder, confirm the roles and reporting channels, complete the Approved AI Tools Register, and align the Green/Yellow/Red guide with your existing classification scheme. Remove this template guidance page or retain it as an administrator cover sheet. Obtain counsel review and executive approval, then distribute the policy and collect acknowledgments.

No product, subscription, or feature is approved merely because it is mentioned in this sample. Approval must identify the exact service, plan, account environment, and permitted use.

1 Purpose and scope

[Organization Name] (the Organization) permits AI only for approved business purposes under this policy. It applies to employees, contractors, consultants, temporary workers, and authorized third parties using AI for Organization work, including on personal devices or while working remotely.

AI includes chatbots, content and image generators, coding assistants, meeting assistants, analytical models, and AI features embedded in ordinary software. It also includes local models, browser extensions, APIs, connected applications, and agents that can access information or perform actions. AI used through a customer or supplier system remains subject to this policy and the applicable contract.

2 Governance and responsibilities

[Executive Sponsor] approves this policy and accepts documented residual risk. [Policy Owner] maintains the policy and register. [IT/Security] evaluates tools, manages access and security controls, investigates incidents, and can suspend unsafe use. [Legal/Privacy/HR] reviews relevant legal, privacy, employment, and contractual issues.

Managers and data owners confirm the business purpose, authorize access to their data, and ensure appropriate human review. Every user must check approval before use, follow data limits, verify outputs, complete training, and promptly report concerns. An AI provider or tool does not assume the user's responsibility for work performed.

[Policy Owner] reviews approvals at least [quarterly] and the policy at least annually. Material changes to vendor terms, model behavior, data flows, permissions, or business use require reassessment before expanded use. Owners must remove access and revoke connections when staff leave, roles change, or approval expires.

3 Approved tools and accounts

Use only a tool and use case listed as approved in Appendix A or the current internal register. Requests go to [AI Approval Channel] and must describe the task, users, data classification, integrations, anticipated benefits, and risks. A trial, free plan, purchased license, or familiar brand is not approval.

Use Organization-managed accounts with the approved business plan, tenant, and configuration. IT must apply single sign-on and multifactor authentication where supported, appropriate access controls, and approved sharing settings. A work email address alone does not make an account Organization-managed.

Personal or unmanaged AI accounts must not receive nonpublic Organization, customer, employee, or partner information. Work using public information through such an account also requires explicit written approval from [IT/Security Approver]; the approval must define the purpose and controls. Do not use personal accounts to evade an Organization restriction.

Before approval, IT/Security and relevant reviewers must assess vendor terms, data use for training, retention and deletion, subprocessors, processing locations, access and audit capabilities, incident support, and contractual protections. Record the approved settings and limits. Paid plans and training opt-outs alone do not establish that a service is safe for confidential data.

4 Data handling and intellectual property

Classify information before entering a prompt, uploading a file, enabling a connector, or recording a meeting. Appendix B defines Green, Yellow, and Red data. If classification is unclear, stop and ask [Data Owner or Security Contact]. Approval of a tool does not authorize every kind of data.

- Use the smallest amount of information needed. Prefer public, synthetic, or appropriately de-identified data. Removing a name alone may leave a person or customer identifiable.
- Protect the Organization's and customers' intellectual property, including source code, designs, business processes, pricing, and nonpublic documents. Obtain data-owner permission and honor customer contracts and confidentiality obligations.
- Do not enter passwords, private keys, access tokens, recovery codes, or other authentication secrets into prompts, uploads, or model memory. Use approved secret-management mechanisms for authorized integrations.
- Use approved storage, sharing, and retention controls for prompts, outputs, files, transcripts, and saved memory. Check links and recipient access before sharing. AI output inherits the sensitivity of the data it contains or reveals.
- Training, fine-tuning, knowledge bases, and bulk ingestion require separate approval covering data rights, quality, representativeness, access permissions, retention, and testing for unintended disclosure. Do not repurpose data beyond its authorized use.

Keep required business records in approved repositories under [Records Retention Schedule]. Delete other AI content according to approved schedules, subject to legal holds. Do not assume deleting a chat removes every vendor copy, backup, connected file, or audit record.

5 Acceptable and prohibited uses

Within a tool's approved limits, staff may use AI to brainstorm, draft or edit text, summarize authorized materials, develop learning aids, analyze permitted data, or assist with code. Users must review the result before relying on it and apply the same confidentiality and quality standards as for other work.

The following activities are prohibited:

- Using unapproved AI with nonpublic data; bypassing security controls; connecting unapproved accounts, extensions, or services; or sharing access credentials.
- Generating deceptive impersonations, harassment, unlawful discrimination, fabricated evidence, or content that violates intellectual property rights, contracts, or Organization policy.
- Allowing AI alone to make consequential employment, credit, legal, medical, financial, safety, or similar decisions. These uses require separate written review by the accountable business owner and relevant specialists, with meaningful human oversight.
- Publishing, sending, deploying, or executing AI output without the review and approvals required for that activity, or presenting unverified output as established fact.

6 Shadow AI and embedded features

Shadow AI is AI used for work without the Organization's approval or visibility. This includes personal chatbot accounts, unapproved meeting bots, browser extensions, local tools, and AI features enabled inside otherwise approved software. Employees must disclose existing business use through [AI Inventory Channel], including tools obtained without an IT purchase.

Approval of an application does not automatically approve its new AI features. Before enabling an embedded assistant, search or summarization feature, plug-in, or AI add-on, check the register or ask IT. Review what it can read, what it sends externally, whether it retains data, and who can access its outputs. Disable unapproved features and stop nonpublic data transfers until reviewed.

IT/Security may use application inventories, vendor administration settings, endpoint controls, network or cloud activity, and employee reports to identify shadow AI. Managers must help staff obtain approved alternatives and report unsafe use. Employees must not conceal use or move it to personal devices to avoid controls.

7 Agents and connected access

An AI agent can take actions, such as searching files, sending messages, modifying records, or running code. A connector or OAuth authorization can grant an AI service access to email, calendars, cloud storage, business applications, or other systems. Treat these grants as access decisions, even when they are presented as a simple sign-in or consent screen.

- IT/Security must approve each agent, connector, OAuth permission grant, API integration, and service identity before use. Staff must not grant organization-wide access or administrative consent.
- Grant only the permissions, data sources, users, and duration needed. Prefer read-only access when it satisfies the task. Protect integration secrets in approved stores; use individual or scoped service identities rather than shared accounts.
- Test in an isolated environment using synthetic or authorized test data. Check for unintended disclosure, excessive permissions, unsafe actions, and prompt injection, where instructions in an untrusted file, website, or message try to redirect the AI.
- Require a designated person to approve external messages, purchases, payments, deletions, production changes, access changes, and other material commitments before execution, unless a separately approved automation defines bounded actions and documented controls.
- Record the owner, access scope, action limits, logs, review date, and method for immediately disabling the agent or revoking tokens. Set appropriate execution and spending limits, and test recovery for changes the agent can make.

Treat retrieved content as untrusted input. Instructions found in a document, web page, or email cannot override this policy or authorize broader access. Pause and report unexpected permission requests, destinations, or actions. Reassess agents whenever their capabilities or connections change.

8 Meetings recordings and transcription

Use only approved recording, transcription, translation, and meeting-summary services. The meeting organizer must confirm the service and data class are approved, provide required notice, and obtain consent where applicable before recording or allowing a bot to join. [Legal/Privacy Contact] determines requirements for participants' locations, meeting type, and customer agreements.

Do not enable automatic recording, silent transcription, or automatic bot attendance by default. Confirm that participants know when recording or transcription is active and how the records will be used. If required permission is absent or withdrawn, stop and follow [Alternative Meeting Procedure].

Exclude sensitive topics unless the meeting and service have explicit approval for that data. Limit transcript, recording, and summary access to authorized recipients; check automatic distribution and sharing links. Apply [Meeting Records Retention Schedule], legal holds, and approved deletion settings. Verify names, decisions, commitments, and action items before distributing AI-generated minutes.

9 Human verification and responsible use

AI can produce plausible but incorrect information, invented citations, insecure code, or biased recommendations. The employee using the output remains responsible for its quality and must involve an appropriately qualified reviewer when the subject or impact requires it.

- Check facts, calculations, source references, and quoted material against reliable original sources. Do not treat the AI's own explanation or another AI answer as sufficient verification.
- Review for confidential information, bias, discriminatory effects, accessibility, intellectual property concerns, and suitability for the intended audience. Use representative test data where appropriate and document material limitations.
- For code, formulas, scripts, or configuration changes, complete the required security review and testing before execution or deployment. Keep AI-generated code out of production until normal change controls are satisfied.
- Clearly disclose material AI involvement when required by law, contract, customer expectations, or [Disclosure Standard]. Do not claim that a person created, reviewed, or endorsed work when that is untrue.

Keep enough documentation to support significant decisions, including relevant sources, review, and approval, without retaining unnecessary sensitive data. People must be able to question consequential AI-assisted decisions through [Review or Appeal Channel]. Explain the basis and limitations in terms the affected audience can understand.

Choose tools and workflows appropriate to the task, considering reliability, accessibility, cost, and environmental impact where practical. Use a simpler approved method when AI adds little value. Raise ethical, safety, or privacy concerns with [Policy Owner or Compliance Contact].

10 Incident reporting and response

Report suspected exposure, unauthorized use, unsafe agent activity, incorrect external disclosure, or unusual AI behavior immediately to [Security Contact and Urgent Channel]. Report near misses as well as confirmed incidents. Do not delay reporting while determining whether harm occurred.

Stop the affected activity. If safe, pause the agent or disconnect the session; ask IT to revoke access and tokens. Preserve relevant messages, timestamps, tool and account details, permissions, and logs under incident-response instructions. Avoid copying exposed sensitive data into a new ticket or AI service, and do not destroy evidence. Do not investigate beyond your authority or contact affected customers, regulators, or vendors on your own.

[Incident Response Lead] coordinates containment, investigation, recovery, and any required notification with counsel and other responsible teams. The Organization documents corrective actions and reassesses affected approvals. Good-faith reporting is encouraged and must not result in retaliation; intentional misuse remains subject to review.

11 Monitoring compliance and exceptions

Subject to applicable law and required notice, the Organization may review business AI activity on its systems and managed accounts to protect information and enforce this policy. Monitoring may include tool use, account and connector access, security alerts, prompts, outputs, and logs where available. Coverage varies by service and configuration; monitoring does not replace employee reporting.

[Monitoring Owner] restricts access to monitoring records, applies [Monitoring Retention Period], and follows [Employee Privacy and Monitoring Notice]. Monitoring must be proportionate to its purpose. Counsel and HR must review relevant notice, privacy, and employment requirements before implementation.

Exceptions require written approval from [Exception Authority], with IT/Security, the data owner, and Legal/Privacy review as needed. Document the reason, data and activities allowed, safeguards, owner, expiry date, and review plan. Exceptions cannot override law, contracts, or the prohibition on entering authentication secrets. Violations may lead to access suspension or disciplinary action, up to termination, consistent with law, contracts, and HR policy.

12 Training adoption and ongoing review

Complete AI training before access and at least [annually], with additional instruction when tools or risks change. Training must cover the data guide, approved accounts, verification, shadow AI, agent permissions, transcription, and incident reporting. [Training Owner] maintains attendance and acknowledgments.

This policy supplements [Organization Acceptable Use Policy] and related security and privacy policies. Apply the stricter applicable requirement and ask [Policy Owner] to resolve conflicts. [Legal/Compliance Owner] must identify obligations based on the Organization's industry, jurisdictions, data, and contracts, including export controls or regulated information where applicable. AI use is permitted only within those obligations and documented approvals.

Review the policy at least annually and after material incidents, legal changes, or major technology changes. [Approving Authority] approves revisions; [Policy Owner] communicates changes and updates the register. Questions go to [Policy Contact].

Employee acknowledgment

I acknowledge that I have received and read the [Organization Name] AI Acceptable Use and Data Security Policy, version [Version], effective [Effective Date]. I understand that I may use AI for work only within approved tool, account, data, and activity limits.

I agree to protect Organization and customer information, verify AI-assisted work, complete required training, and promptly report suspected incidents or unauthorized use. I understand that exceptions require written approval and that violations may result in action under applicable Organization policies. I know how to contact [Policy Contact] if I need clarification.

Employee or worker name: [Full name]

Role or department: [Role or department]

Signature: _____

Date: _____

Policy version acknowledged: _____

Training completed on: _____

Return this acknowledgment to [HR or Policy Owner / Submission Method]. An approved electronic acknowledgment may be used under the Organization's procedures.

Administrator adoption checklist

Complete before issuing the policy. Retain this checklist with the adoption record or remove it from the employee copy.

- All placeholders completed; policy contacts and reporting routes tested.
- Approved tool entries and supporting approval records completed and published.
- Data guide aligned with existing classification, customer restrictions, and regulated data.
- Account, connector, recording, retention, and monitoring controls confirmed.
- Counsel, security, privacy, and HR reviews completed; executive approval recorded.
- Employee training scheduled and acknowledgment collection assigned.

Approved by: [Name and title] | Approval date: [Date]

Counsel review completed by: [Reviewer] | Review date: [Date]

Appendix A Approved AI Tools Register

This register is maintained by [Register Owner] at [Internal Location]. Blank entries and pending reviews do not authorize use. Each approval applies only to the named service, plan, environment, features, and use case. Duplicate rows and approval records as needed.

Tool and account	Permitted use and data	Controls and connections	Owner and status
[Tool / vendor] [Plan / tenant] [Managed account type]	[Approved use cases] [Green only or specified Yellow] [Prohibited activities]	[Approved features] [Settings / record link] [Permitted connectors]	[Business owner] [Approver / approval date] [Status / next review]
[Tool / vendor] [Plan / tenant] [Managed account type]	[Approved use cases] [Green only or specified Yellow] [Prohibited activities]	[Approved features] [Settings / record link] [Permitted connectors]	[Business owner] [Approver / approval date] [Status / next review]

Supporting approval record for each entry

Tool and use case: [Name, service URL, exact product or feature, plan, tenant, version where relevant, intended task, user group, and business justification]

Data authorization: [Permitted classification and specific data categories, data owner, excluded data, contractual restrictions, and approved processing locations]

Vendor review: [Terms and privacy review date, agreement or assessment reference, subprocessors, model-training use and opt-out settings, retention and deletion limits, and incident support contact]

Identity and security: [SSO/MFA settings, roles, sharing restrictions, approved devices, audit capabilities, monitoring coverage, and any control gaps]

Agent and connector scope: [Approved data sources, OAuth permissions, service identity, allowed actions, required human approvals, execution/spending limits, disable or token-revocation procedure; enter None if not applicable]

Records and testing: [Prompt/output/transcript retention, record location, legal-hold process, verification requirements, test evidence, known limitations, and recovery plan]

Approval and lifecycle: [IT/Security approval, business/data-owner approval, Legal/Privacy review as needed, residual-risk acceptance, approval/expiry dates, next review, and reassessment triggers]

Current status: [Approved / Approved with restrictions / Pending / Suspended / Retired]. Any restrictions and exceptions: [Details, approver, expiry, and record reference].

Vendor claims must be verified against the applicable plan, contract, and actual configuration. Re-review the entry when permissions, model behavior, terms, retention, data use, or enabled features materially change.

Appendix B Green Yellow Red data guide

Apply this guide before prompting, uploading, connecting, recording, or enabling model memory. The most sensitive information in a file or conversation determines its classification. Colors supplement the written labels; approval and contractual limits always apply.

Class	Examples	Required action
GREEN Public or synthetic	Published website copy; publicly released product information; genuinely synthetic examples that reveal no real person, customer, secret, or nonpublic operation.	Use an approved tool and account for an approved task. Check rights and accuracy. Public availability does not remove copyright or contractual restrictions.
YELLOW Nonpublic business data	Internal procedures; nonpublic pricing; business correspondence; customer documents; proprietary code; ordinary personal information where approved. Some of these may be Red under local rules.	Use only the exact managed service approved for that data and purpose. Obtain data-owner authorization, apply approved vendor/retention controls, and minimize or de-identify information. Stop if approval is unclear.
RED Restricted or highly sensitive	Authentication secrets; highly sensitive employee or customer records; health or financial account data; legally privileged material; export-controlled information; sensitive security details.	Do not enter into ordinary AI services. Non-secret Red data requires a separately approved restricted environment and documented Legal/Privacy, Security, and data-owner review. Never put authentication secrets in prompts or uploads.

Everyday examples

- Drafting a public event invitation using published details: Green, using an approved tool.
- Summarizing a customer proposal or pasting internal code: Yellow or Red depending on its contents; confirm data-specific approval first.
- A meeting that discusses personnel issues, health information, legal advice, or account credentials: treat the sensitive content as Red; do not admit an ordinary transcription bot.
- Connecting AI to a mailbox or shared drive: classify all information it could access, not only the document you intend to summarize. Approve and restrict the connection before use.

If unsure, do not submit the information. Ask [Data Owner or Security Contact]. Redacting names, using a paid account, or turning off model training does not automatically lower the classification or authorize use.

Administrator customization and reference

Map these labels to [Organization Classification Standard]. Identify industry, customer, and jurisdiction-specific requirements with counsel. Examples are illustrative and must be adjusted to the Organization's obligations; this guide is not a compliance determination.

Optional governance reference: NIST AI Risk Management Framework and Generative AI Profile, NIST AI 600-1. These provide voluntary risk-management guidance, not a substitute for legal review. <https://www.nist.gov/itl/ai-risk-management-framework>

Brand wording reference: Just Solutions, Inc., How to Use ChatGPT for Your Business. <https://justsolutions.com/blog/how-to-use-chatgpt-for-your-business-beginner-friendly-guide/>