withum

Sample "Acceptable Employee AI Use" Policy

It is recommended to review the template with your General Counsel before implementation.

Purpose: This policy defines how employees may use Artificial Intelligence (AI) tools in the course of work for **[Organization Name]**. It is intended to enable responsible innovation while protecting our clients, employees, and the organization.

Scope: *This policy applies to*

- All employees, contractors, and temporary workers ("employees")
- All AI systems used for work purposes, including public tools, vendor-embedded AI features, and internal AI solutions
- All work locations and devices, including remote work

Definitions (customize as needed):

- **AI tools / systems:** Technologies that generate, transform, classify, recommend, or summarize content using machine learning or generative models (e.g., LLMs).
- **Public AI tools:** Externally hosted tools not specifically contracted/approved by the organization.
- **Approved AI tools:** Tools reviewed and authorized by **[AI Governance Owner / Group]**.
- **Sensitive data:** Any information classified as **[Confidential / Restricted / Regulated]**, including client data, employee data, PHI, PCI, IP, credentials, and non-public business information.
- **Personal data:** Data relating to an identifiable individual as defined by applicable privacy laws.
- **High-risk use case:** A use involving material decisions, regulated data, or external impact *(see Section 3)*.

Policy owner:	[Title/Team]
Effective date:	[Date]
Review cadence:	At least **[annually / semi-annually]** or upon major regulatory/technology changes.

SECTION 1: RISK APPETITE AND AI USAGE PRINCIPLES

1. Risk appetite statement (choose/edit)

Example — Innovation-forward posture

"[Organization Name] is an innovation-first organization. We encourage employees to explore AI tools that improve quality, speed, or learning, provided usage complies with this policy, client obligations, and data protection requirements."

Example — Cautious posture

"[Organization Name] values new technology and productivity gains, but we will not accept unnecessary risk with unvetted AI. Employees must use only approved AI tools for work, and must submit any new AI use case for review before adoption."

2. AI usage principles (template)

Employees must follow these principles whenever using AI for work:

1	**Protect data** + Do not input sensitive or personal data into unapproved tools. + Use minimum-necessary data and prefer anonymized or synthetic data.
2	**Maintain human accountability** + Employees remain responsible for all outputs, decisions, and deliverables. + AI is an assistive tool, not an authority.
3	**Be transparent** + Disclose AI involvement to stakeholders where relevant, especially for client-facing deliverables or material internal decisions. + Follow any contract/client rules about disclosure.
4	**Ensure fairness and non-discrimination** + Do not use AI to make or automate decisions that may create bias without review and controls. + Avoid using demographic data or proxies unless explicitly approved.
5	**Respect law, policy, and ethics** + Follow applicable regulations (privacy, IP, employment, sector rules). + Do not use AI to mislead, impersonate, or manipulate.
6	**Verify quality** + Treat AI outputs as drafts. + Validate facts, citations, calculations, and reasoning prior to use.

Optional: link to your Responsible AI standard / code of conduct **[insert link]**.

SECTION 2: AI TECHNOLOGIES EMPLOYEES CAN USE

1. Approved tools list (template)

Employees may use only the tools listed below for work purposes unless granted a documented exception.

Tool / Product	Approved Use	Data Allowed	Notes / Restrictions
[Tool A]	[e.g., drafting, summarization]	[e.g., internal public data only]	[e.g., enterprise instance; no training on our data]
[Tool B]	[..........................]	[..........................]	[..........................]

Owner of approved list: [AI Governance Owner / IT Security]

Where list is published: [intranet link / wiki]

2. Public and unapproved tools

Unless explicitly approved:

- Do not use public AI tools for work involving any sensitive, confidential, or personal data.
- Do not connect unapproved AI tools to company systems (email, storage, CRM, ticketing, etc.).
- Do not install unapproved AI plugins/extensions in browser or software.

3. Device and environment requirements

- Business AI use must occur only on **[company-managed / approved]** devices and accounts.
- Using personal devices for business AI work is **[prohibited / allowed only with MDM/encryption]**.
- Employees must not bypass security controls (e.g., using personal logins to access prohibited features).

4. Vendor AI features embedded in tools

Some approved enterprise platforms may include AI (e.g., meeting summary, email drafting). Employees may use these features only if the platform itself is approved and the feature is enabled by **[IT/AI Governance]**.

SECTION 3: ACCEPTABLE AI USES

This section defines ***Always Allowed, Always Prohibited, and Allowed Only With Review*** uses.

1. Always allowed uses (example set)

Employees may use approved AI tools for the following low-risk activities without pre-approval, as long as no sensitive/personal data is used:

+ Drafting or editing internal content (emails, outlines, slide text)
+ Summarizing meeting notes or internal documents
+ Brainstorming ideas, alternatives, or plans
+ Translating non-confidential text
+ Writing or refactoring code **using non-sensitive inputs** and with human review
+ Creating synthetic examples to test concepts

Key rule: outputs must be reviewed for accuracy, tone, legality, and client impact.

2. Always prohibited uses

Employees must never use AI for the following:

Data and privacy

+ Inputting or uploading sensitive or personal data into unapproved/public tools
+ Using AI to re-identify anonymized individuals
+ Training models on protected data without documented consent/approval
+ Copying client confidential info into AI prompts unless tool and use are approved

Security

+ Sharing credentials, API keys, secrets, or system architecture in AI prompts
+ Using AI to develop or distribute malware, exploit code, or unauthorized access methods
+ Bypassing company security or monitoring controls

Decisioning and people impacts

+ Fully automated hiring, promotion, termination, compensation, or performance decisions
+ Using AI to score, rank, or profile individuals without review and a validated fairness plan
+ Surveillance or monitoring of employees/clients beyond approved business processes

External misuse

+ Generating deceptive content (impersonation, falsified evidence, undisclosed synthetic media)
+ Publishing AI outputs externally as factual without verification
+ Using AI in ways that violate law, contracts, or professional standards

3. Uses requiring review before proceeding

Employees must submit a use case for review if it involves any of the following:

- **Material decisions** affecting employees, clients, finances, eligibility, or safety
- **Client-facing or public-facing** AI outputs, deliverables, or interactions
- Any use of **confidential, regulated, or personal data**
- Integrating AI into workflows that run automatically (bots, agents, pipelines)
- Building or fine-tuning models using organizational or client data
- Any AI use in regulated domains: **[e.g., healthcare, finance, government]**

If unsure, assume review is required.

4. Content and IP expectations

- Employees must ensure AI outputs do not infringe copyrights, licenses, or third-party IP.
- Do not paste proprietary third-party content into tools unless allowed by license.
- Cite sources for factual claims when relevant, and do not fabricate citations.

SECTION 4: PROCESS FOR REQUESTING USE CASE REVIEW

1. When to submit

Submit a review request **before** starting work if your use case fits Section 3.3 or you're uncertain.

2. How to submit (template)

Employees should submit requests via ***[form/link/email]*** *with:*

- **Use case title and owner**
- **Business purpose and expected benefits**
- **AI tool(s) requested** and whether tool is already approved
- **Data involved** (types, classification, volume, source)
- **Outputs and who will use them** (internal/external)
- **Potential risks** (privacy, bias, security, legal, reputational)
- **Safeguards** (human review, logging, red-teaming, bias testing)
- **Go-live** timeline and required support

3. Review participants (template)

Requests will be evaluated by ***[AI Governance Committee]****, typically including:*

- Business sponsor(s)
- Privacy / data protection lead
- Legal / compliance lead
- Cybersecurity / IT lead
- AI/ML or analytics expert(s)
- HR lead (if people impacts)

4. Outcomes

Possible outcomes include:

- **Approved** (no conditions)
- **Approved with conditions** (e.g., tool restrictions, monitoring, limited data)
- **Pilot only** (time-boxed test with evaluation gates)
- **Not approved** (with rationale)

Employees must follow any conditions listed in the approval record.

SECTION 5: AI USAGE AND RISK MONITORING

1. Monitoring approach (template)

To maintain safety and compliance, ***[Organization Name]*** *may:*

- Log usage of approved AI tools (metadata, not necessarily content)
- Review high-risk use cases periodically for performance, drift, and bias
- Audit access patterns, integrations, or data exposure pathways
- Require re-approval after major tool/model changes

Monitoring will be performed consistent with **[privacy policy / employee monitoring policy]**.

2. Employee responsibilities

Employees using AI must:

- Follow approved use case conditions
- Maintain records of AI involvement where required
- Immediately report suspected misuse, data leakage, or harmful outcomes
- Cooperate with audits or reviews
- Stop using a tool or workflow if instructed by **[AI Governance / Security]**

3. Incident reporting

Report AI-related incidents via ***[security/privacy hotline or ticket link]*** *within* ***[X hours/days]****, including:*

- Tool used
- Data involved
- What happened and when
- Known impact and who may be affected
- Steps already taken

Acknowledgement

All employees must acknowledge they have read and will comply with this policy, and complete required training **[annual / onboarding]**.

Possible outcomes include:

- Approved (no conditions)
- Approved with conditions (e.g., tool restrictions, monitoring, limited data)
- Pilot only (time-boxed test with evaluation gates)
- Not approved (with rationale)

Employees must follow any conditions listed in the approval record.

Employee name: ____

Date: _________________________

withum

Appendix A: AI Use Policy FAQ + Practical Examples

This appendix provides plain-English guidance and examples to help employees apply the AI Acceptable Use Policy in day-to-day work. If you are unsure whether a use is allowed, pause and ask for review per Section 4.

FAQ

1) Why do we have an AI use policy?

Because AI can be extremely helpful — and risky if used incorrectly.

This policy enables responsible innovation while protecting our clients, employees, and the organization's data, reputation, and legal obligations.

2) Can I use AI at work at all?

Yes. Employees are encouraged to use approved AI tools for approved low-risk tasks (see section 3.1) such as drafting, summarizing, brainstorming, and internal productivity support — as long as you don't share sensitive or personal data.

3) What counts as "sensitive" or "personal" data?

Use your organization's data classification standard, but examples normally include:

- Client confidential information
- Non-public financials, pricing, contracts, deal terms
- Employee records, compensation, performance data
- Health data (PHI), payment data (PCI), tax IDs, SSNs
- Credentials, API keys, secrets, internal system details
- Anything labeled [Confidential / Restricted / Regulated]

If you wouldn't paste it into a public website or email to the wrong person, don't put it into AI unless your tool and use case are explicitly approved.

4) What's the difference between an "approved tool" and a "public tool"?

- **Approved tool:** reviewed and authorized for internal use by **[AI Governance Owner / IT Security]**, often under an enterprise contract.
- **Public/unapproved tool:** anything not on the approved list, including personal accounts or free web versions.

Rule: Use only approved tools for business. Never upload sensitive/personal data to public tools.

5) Can I paste client material into an approved AI tool?

Only if BOTH are true:

- The tool is approved for that data type, and
- Your use case involving client data is approved (Section 3.3).

When in doubt, submit a review request.

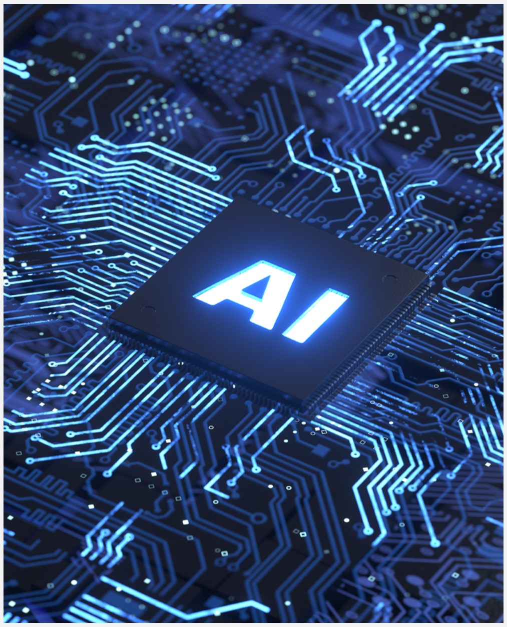

6) Can I use AI to write client deliverables?

Maybe — but usually only after review.

Even if the tool is approved, externally facing work often requires review so we confirm accuracy, disclosure obligations, and IP risk.

7) Do I need to tell anyone I used AI?

Sometimes.

- For internal low-risk work, disclosure is optional unless your department requires it.
- For client-facing deliverables, public content, or material decisions, disclose AI involvement per Section 1.2 (Transparency).

8) Who is responsible for the output if AI makes a mistake?

You are.

AI is an assistive tool. Employees must validate outputs before use (Section 1.2, Quality + Accountability).

9) Can I use AI to help with hiring or performance reviews?

Not without formal approval.

Anything influencing employment decisions is a high-risk use case and must be reviewed.
Fully automated decisioning is prohibited.

10) Can I use AI coding assistants?

Yes, with caution and only in approved environments.

+ Don't paste secrets or proprietary client code into unapproved tools
+ Validate security and licensing.
+ You remain responsible for correctness.

11) What if AI gives me something that looks confident but wrong?

Treat AI outputs as drafts. Verify facts, sources, calculations, and logic before using or sharing.

12) What if I'm not sure whether something is allowed?

Assume it requires review.

Submit a request via [intake link]. Quick questions can go to [alias/Slack channel].

13) What happens if someone violates this policy?

Violations may result in revocation of AI access, disciplinary action, and/or legal response depending on severity — especially where client data or security is involved.

EXAMPLES BY CATEGORY

A) Always Allowed (low-risk, using approved tools, no sensitive/personal data)

1. Summarizing internal meeting notes

+ **OK:** Summarize today's project standup into 5 bullets.
+ **Why:** Low impact, no sensitive data.

2. Drafting internal emails

+ **OK:** Rewrite this note to be clearer and more professional.
+ **Why:** Productivity support; human remains accountable.

3. Brainstorming

+ **OK:** Give me 10 ways to structure a client workshop on cloud cost optimization.
+ **Why:** Idea generation; no restricted data.

4. Improving clarity of non-confidential documents

+ **OK:** Edit this internal FAQ for tone and readability.
+ **Why:** Content polishing; low risk.

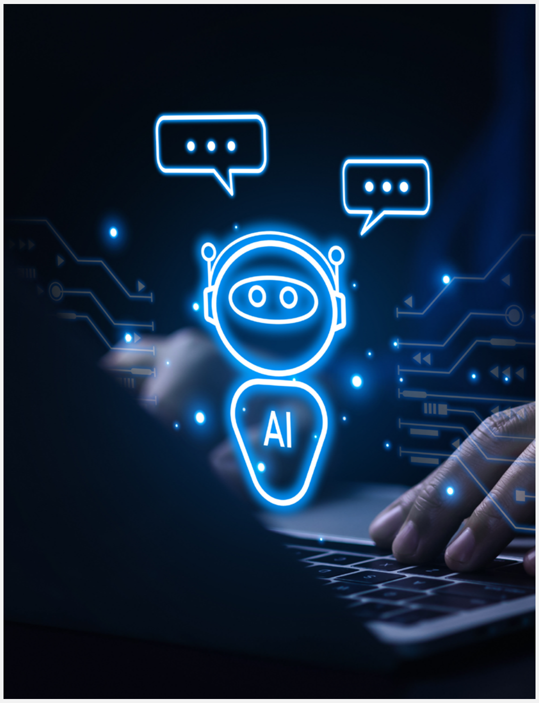

5. Creating synthetic examples

+ **OK:** Create a mock dataset for testing a dashboard; use fake names.
+ **Why:** Synthetic data avoids privacy risk.

B) Always Prohibited

1. Uploading sensitive or personal data into public tools

+ **Not OK:** Here's a client contract — summarize key obligations. (using public ChatGPT)
+ **Why:** Client confidential data in an unapproved tool.

2. Sharing credentials or secrets

+ **Not OK:** Here's our AWS key — help me debug why it fails.
+ **Why:** Security risk + policy violation.

3. Fully automated people decisions

+ **Not OK:** Rank these candidates automatically and decide who to hire.
+ **Why:** High-risk employment decisioning; prohibited.

4. Creating deceptive or impersonation content

+ **Not OK:** Write an email pretending to be the client to trick a vendor.
+ **Why:** Ethical/behavioral violation.

5. Generating malware/exploit code

+ **Not OK:** Write a script to break into a system.
+ **Why:** Security and legal breach.

C) Requires Review Before Proceeding

1. Client-facing deliverables

+ **Review needed:** Generate a first draft of the client strategy deck.
+ **Why:** External impact, disclosure, IP, accuracy risks.

2. Using confidential company data

+ **Review needed:** Analyze internal pricing margins and propose changes.
+ **Why:** Confidential data used in AI workflow.

3. Any regulated/personal data

+ **Review needed:** Summarize employee engagement survey comments containing names.
+ **Why:** Personal data.

4. Bots/agents integrated into systems

- **Review needed:** Deploy a Slack bot that auto-responds to client tickets.
- **Why:** Automated workflow + external interaction.

5. AI influencing financial or legal outcomes

- **Review needed:** Use AI to recommend contract redlines.
- **Why:** Material legal impact; risk of hallucination.

Example Prompts: Safe vs Unsafe

Safe prompt pattern	Unsafe prompt pattern
"Here is a non-confidential summary of the situation. Please draft options for me."	"Here are client details / employee names / financials — analyze and recommend actions."
"Assume all names are fictional. Generate a sample for testing."	"Use our internal strategy doc verbatim to draft a public blog."
"Do not invent facts; if uncertain, say so."	"Generate citations even if you don't know the source."

"If This, Then That" Quick Guide

- **If you're using an approved tool and no sensitive/personal data:** usually ok.
- **If you're using a public/unapproved tool:** only OK for generic work with zero confidential data.
- **If it affects clients, employees, money, or reputation:** submit for review.
- **If you're unsure:** submit for review.

Optional Add-On: Short Disclosure Language (Template)

Use this when disclosure is required:

Internal	Client/public (if permitted)
"AI was used to assist drafting/summarization. Final content was reviewed by **[name/role]**."	"Portions of this deliverable were generated with AI assistance and validated by **[Organization Name]** subject-matter experts."