AI Tool Security Vetting Kit

20-Question Checklist · Data Classification Matrix

Sample DPA Clauses · NIST AI RMF Mapping · SOC 2 + EU AI Act Crosswalk · Approve / Approve-with-Controls / Reject Decision Tree

---

85%

of employees use at least one AI tool not approved by IT (Gartner 2025)

WHAT IS IN THIS KIT?

PART 1 OF 2 — QUESTIONS 1–10

20-Question AI Vendor Vetting Checklist

Run through these 20 questions before approving any AI tool for company use. Any "No" or "Unknown" answer requires a documented risk exception or vendor remediation.

1 DATA HANDLING

Does the vendor publish a clear data processing agreement (DPA) or AI-specific addendum?

Why: Without a DPA, you have no contractual control over what happens to data you submit.

2 TRAINING DATA EXCLUSION

Does the vendor contractually commit to NOT using your data to train their models?

Why: Many default API tiers use prompts/responses for model improvement unless you opt out or upgrade to an Enterprise plan.

3 DATA RETENTION

What is the vendor's data retention period for prompts and outputs? Can you request deletion?

Why: Default retention of 30–90 days for prompt history is common. Regulated data should not persist beyond the session.

4 SUB-PROCESSORS

Does the vendor publish a current sub-processor list? Are any sub-processors in jurisdictions with weak data protection laws?

Why: Sub-processors inherit your data. A missing or outdated list is a red flag for a SOC 2 CC9.2 audit.

5 DATA RESIDENCY

Where is your data processed and stored? Is residency in your required region (EU, US, UK) available?

Why: GDPR Art. 44–49 restrict transfers outside the EEA without adequate safeguards. Financial regulators often require US processing.

6 ENCRYPTION

Is data encrypted in transit (TLS 1.2+) and at rest (AES-256 or equivalent)?

Why: Encryption is table-stakes for any SaaS vendor receiving confidential data. Absence should be disqualifying.

7 AUDIT RIGHTS

Does the contract include audit rights, SOC 2 Type II report sharing, or ISO 27001 certification?

Why: SOC 2 CC9.2 requires documented vendor oversight. An AI vendor with no audit evidence cannot be Tier 1 approved.

8 INCIDENT NOTIFICATION

What is the vendor's breach notification timeline? Does it meet your shortest regulatory deadline (72 hours for GDPR/HIPAA)?

Why: A vendor with a 30-day breach notification window makes GDPR 72-hour reporting impossible for any incident touching their platform.

9

ACCESS CONTROLS

Does the vendor support SSO/SAML integration for enterprise accounts? Is MFA enforced for admin access?

Why: Individual employee accounts without SSO create shadow IT — impossible to revoke centrally when someone leaves.

10

HALLUCINATION & OUTPUT RISK

Does the vendor document known accuracy limitations? Have you tested outputs against your highest-stakes use cases?

Why: AI hallucinations in legal, medical, financial, or compliance contexts create liability. Vendor documentation of limitations is a baseline due diligence requirement.

PART 2 OF 2 — QUESTIONS 11–20

20-Question AI Vendor Vetting Checklist (continued)

11 IP OWNERSHIP

Who owns the outputs generated by your prompts? Does the vendor claim any license over your generated content?

Why: Most AI vendors grant you ownership of outputs but retain a license to use them for service improvement. Enterprise plans typically remove this license.

12 CONFIDENTIALITY

Does the DPA include explicit confidentiality obligations for your data? Does it extend to sub-processors?

Why: Confidentiality clauses must flow down to sub-processors. A DPA that doesn't bind sub-processors provides limited protection.

13 REGULATORY COMPLIANCE

Has the vendor completed a HIPAA BAA (if applicable), GDPR Art. 28 DPA, or other sector-specific compliance agreement?

Why: Processing PHI through an AI tool without a BAA is an immediate HIPAA violation regardless of the breach outcome.

14 VULNERABILITY MANAGEMENT

Does the vendor publish a vulnerability disclosure program or bug bounty? What is their patch SLA?

Why: A vendor with no published vulnerability management process cannot demonstrate NIST AI RMF MANAGE compliance.

15 MODEL PROVENANCE

Is it clear which underlying model(s) are used? Have you assessed whether those base models were trained on ethically sourced data?

Why: Regulatory scrutiny on AI training data (EU AI Act Art. 10, FTC guidance) is accelerating. Knowing your model lineage is early-stage due diligence.

16 BIAS & FAIRNESS

Has the vendor conducted bias and fairness testing? Are results documented and available for high-risk use cases?

Why: EU AI Act classifies certain AI uses as "high risk" requiring bias assessment. US EEOC and financial regulators are issuing similar guidance.

17 PROMPT INJECTION RISK

If the AI tool processes external content (emails, documents, URLs), has prompt injection risk been assessed and mitigated?

Why: Prompt injection can cause an AI agent to exfiltrate data or take unauthorized actions. Agents with access to internal systems require formal threat modeling.

18 VENDOR CONCENTRATION

If this vendor had an outage or was breached, what is your fallback? Is there dangerous business process concentration on a single AI provider?

Why: The MOVEit and Snowflake breaches demonstrated that concentrating business-critical workflows in a single SaaS creates systemic risk.

19 EMPLOYEE TRAINING

Have employees using this tool been trained on: what data is safe to input, how to verify outputs, and how to report anomalies?

Why: Tool approval without user training is security theater. The Samsung incident happened because employees lacked a data classification policy for AI inputs.

20 REVIEW CADENCE

Is there a documented review date (annual minimum) for this approval? Who is responsible for re-evaluating if the vendor's terms change?

Why: AI vendor terms change frequently. OpenAI changed its enterprise data policy twice in 2023–2024. A one-time approval with no review cycle is not a control.

DATA CLASSIFICATION

What Data is Allowed in Which AI Tool?

Use this matrix before submitting any data to an AI tool. "Approved" means the data type is generally safe for that tier of tool. "Approved with controls" means extra steps are required. "Not permitted" means the data must never go into that tool class.

Public

Information already publicly available: marketing copy, press releases, published research.

ChatGPT/OpenAI	Microsoft Copilot	Claude (Anthropic)	Gemini (Google)
Approved	Approved	Approved	Approved

Internal

Internal business information not public but not regulated: meeting notes, internal policies, non-sensitive project descriptions.

ChatGPT/OpenAI	Microsoft Copilot	Claude (Anthropic)	Gemini (Google)
Approved w/ Controls	Approved	Approved w/ Controls	Approved w/ Controls

Controls required: Requires Enterprise plan (no training opt-out), SSO enforced, DPA signed.

Confidential

Sensitive business data: M&A strategy, unreleased financials, trade secrets, customer lists, pricing strategy, unpublished IP.

ChatGPT/OpenAI	Microsoft Copilot	Claude (Anthropic)	Gemini (Google)
Not Permitted (Consumer/API) Approved w/ Controls (Enterprise)	Approved w/ Controls	Approved w/ Controls	Not Permitted (Consumer) Approved w/ Controls (Workspace)

Controls required: Requires Enterprise/Workspace plan, DPA + confidentiality addendum, no sub-processor in restricted jurisdiction, documented approval.

Regulated

Data subject to legal protection: PHI (HIPAA), PII (GDPR/CCPA), PCI cardholder data, ITAR-controlled technical data, attorney-client privileged information.

ChatGPT/OpenAI	Microsoft Copilot	Claude (Anthropic)	Gemini (Google)
Not Permitted (no ITAR) HIPAA: Approved w/ BAA only	Approved w/ BAA (HIPAA only)	Not Permitted (ITAR) HIPAA: Approved w/ BAA	Not Permitted (ITAR) HIPAA: Approved w/ BAA

Controls required: BAA required for PHI. PCI: zero tolerance — no card data in any AI tool without explicit scoping and QSA sign-off. ITAR: never input without legal review.

DPA CLAUSE LIBRARY — PART 1

Sample DPA + AI Addendum Clauses

Use these clauses as a starting point when negotiating with AI vendors. Clauses marked [REQUIRED] are non-negotiable for regulated environments. Clauses marked [RECOMMENDED] add meaningful protection but may require Enterprise pricing.

C1 REQUIRED

Training Data Exclusion

Vendor shall not use Customer's data, prompts, or outputs to train, fine-tune, or improve Vendor's AI models or any third-party AI model. This restriction applies to all data submitted via the API, Enterprise interface, or any integration.

Vendor status: ChatGPT Enterprise: default for Enterprise API. Consumer API: must opt out in settings. Copilot M365: off by default for enterprise tenants. Claude: configurable via API settings. Gemini Workspace: admin-configurable.

C2 REQUIRED

Data Retention and Deletion

Vendor shall retain Customer data for no more than 30 days from the date of submission, after which all data shall be securely deleted. Customer may request deletion at any time, and Vendor shall confirm deletion within 72 hours.

Vendor status: ChatGPT Enterprise: configurable. Gemini Workspace: 30-day default. Claude API: no persistent storage by default. Always verify Enterprise plan terms — consumer plans may differ.

C3 REQUIRED

Sub-Processor Notification

Vendor shall maintain and make available a current list of all sub-processors that may access Customer data. Vendor shall provide Customer with at least 30 days' prior written notice before adding or replacing any sub-processor. Customer may object to any new sub-processor within 10 business days of notification.

Vendor status: OpenAI sub-processors include Microsoft Azure, Stripe, and others — published at openai.com/policies/sub-processors. All four major vendors publish sub-processor lists for enterprise accounts.

C4 REQUIRED

Breach Notification

Vendor shall notify Customer of any confirmed or suspected security incident involving Customer data within 24 hours of Vendor's discovery. Notification shall include: nature of the incident, categories and volume of data affected, likely consequences, and remediation measures taken or proposed.

Vendor status: GDPR requires 72-hour notification to supervisory authority. A 24-hour vendor notification window gives you time to assess before your regulatory clock starts. Do not accept 30-day or "prompt" notification language.

DPA CLAUSE LIBRARY — PART 2

C5 REQUIRED

IP Ownership and License

All outputs generated by the AI service using Customer data, prompts, or instructions shall be the sole property of Customer. Vendor hereby assigns to Customer all right, title, and interest in such outputs. Vendor retains no license to use Customer outputs for any purpose other than providing the contracted service.

Vendor status: Most enterprise plans include output ownership. Verify the specific Enterprise Order Form or Addendum — the base terms may not include assignment language.

C6 RECOMMENDED

Confidentiality Obligation

Vendor agrees to maintain the confidentiality of all Customer data processed through the service, using at minimum the same standard of care Vendor uses to protect its own confidential information, but in no case less than reasonable care. This obligation survives termination of the agreement for a period of five years.

Vendor status: Standard in most enterprise AI vendor agreements. Confirm the obligation flows to sub-processors — "Vendor shall require sub-processors to maintain equivalent confidentiality obligations."

C7 RECOMMENDED

Audit Rights

Customer or Customer's designated auditor may, upon 30 days' prior written notice, conduct a security audit or review of Vendor's systems and controls relevant to Customer data processing. In lieu of an on-site audit, Vendor may provide a current SOC 2 Type II report, ISO 27001 certificate, or equivalent third-party assessment.

Vendor status: Most vendors satisfy this via annual SOC 2 Type II report. Request the current report before signing — a vendor that cannot provide one within 30 days of your request should not be Tier 1 approved.

C8 REQUIRED (HIPAA)

HIPAA Business Associate Agreement

To the extent Vendor processes Protected Health Information (PHI) as defined under HIPAA, Vendor agrees to execute a Business Associate Agreement (BAA) in a form acceptable to Customer. No PHI shall be submitted to Vendor's service until such BAA is fully executed. Vendor shall report any Breach of Unsecured PHI within 24 hours of discovery.

Vendor status: OpenAI offers a HIPAA BAA for Enterprise accounts. Microsoft Copilot + Azure OpenAI: BAA available through M365 Compliance Center. Google Gemini: BAA available for Workspace Enterprise Plus. Claude (Anthropic): BAA available for qualifying enterprise accounts.

NIST AI RMF (AI 100-1) — GOVERN + MAP

NIST AI RMF Control Mapping

The NIST AI Risk Management Framework (AI RMF 1.0, published January 2023) provides a voluntary framework for managing AI risk. The four functions — GOVERN, MAP, MEASURE, MANAGE — map directly to your AI vendor vetting program.

NIST AI RMF — GOVERN

GV-1.1
AI Risk Policy
Establish and maintain a policy that defines acceptable AI use, prohibited use cases, approval authority, and employee obligations. Map to Q1–Q4, Q9, Q19.

GV-1.2
AI Inventory
Maintain a current inventory of all AI tools in use — including shadow IT. Every tool in the inventory must have a completed vetting checklist (this document) and a documented approval status.

GV-2.1
Roles & Accountability
Assign ownership for AI risk decisions. The CISO or equivalent owns policy; a designated AI Risk Owner approves each tool; each department head is accountable for use within their team.

GV-3.1
Third-Party Obligations
Require all AI vendors to meet your DPA requirements (clauses C1–C8 in this kit). Include AI vendor governance in your annual vendor review cycle.

NIST AI RMF — MAP

MP-1.1
Use Case Risk Classification
For each AI tool, document its intended use case and classify risk: Low (content drafting, search), Medium (customer-facing decisions, internal policy generation), High (HR decisions, credit, medical advice, legal research). High-risk use cases require board or executive sign-off.

MP-2.2
Data Flow Mapping
Map what data enters the AI tool, how it is processed, and where outputs are stored or used downstream. Use the Data Classification Matrix (Page 4) to validate each data type against the tool's approval tier.

MP-3.1
Threat Modeling
Document key threats for each approved AI tool: prompt injection (especially for agentic tools), output manipulation, supply chain risk (vendor breach), and data exfiltration via prompt history.

NIST AI RMF — MEASURE + MANAGE

NIST AI RMF — MEASURE

MS-1.1

Output Accuracy Testing

Before deploying an AI tool for a business-critical workflow, conduct structured accuracy testing with representative samples. Document test methodology, pass/fail criteria, and results. Reassess annually or after major model updates.

MS-2.1

Bias and Fairness Testing

For any AI tool used in decisions affecting employees or customers (hiring, credit, benefits, communications), conduct bias testing across protected categories. Document results. EU AI Act (2024) mandates bias assessment for high-risk systems.

MS-2.5

Vendor Security Assessment

Complete this 20-question checklist (Page 2–3) for each new AI vendor and re-run annually. Score vendors: >85% pass = Tier 1 (Approved), 70–85% = Tier 2 (Approved w/ Controls), <70% = Tier 3 (Reject or defer).

MS-3.3

Monitoring and Anomaly Detection

Implement logging for AI tool use where technically feasible. Monitor for anomalous usage patterns: excessive data submission, use by terminated employees, access outside working hours, or use of unapproved tools at the network layer.

NIST AI RMF — MANAGE

MG-1.1

Incident Response for AI

Extend your IR plan to cover AI-specific incidents: model poisoning, prompt injection attacks, unauthorized training data access, hallucination-driven decisions with material impact. Assign a dedicated responder for AI incidents.

MG-2.2

Unapproved Tool Response

Establish a clear process for when an employee is found using an unapproved AI tool: immediate data assessment (what was submitted?), tool quarantine or access revocation, notification to DPO/CISO, incident logging, and documented remediation.

MG-3.2

Vendor Offboarding

When decommissioning an AI vendor relationship: request written confirmation of data deletion within 30 days, revoke SSO/API credentials, update the AI inventory, and conduct a brief after-action review on what data was exposed during the relationship.

COMPLIANCE CROSSWALK — PART 1

SOC 2 Trust Service Criteria + EU AI Act Crosswalk

The table below maps the most-asked AI vendor questions to the relevant SOC 2 TSC control and EU AI Act provision. Use this when responding to customer questionnaires or internal audit requests.

Control	SOC 2 TSC	EU AI Act	NIST AI RMF
AI vendor inventory maintained	CC9.2 — Vendor risk management	Art. 28 — Obligations of providers (distributor accountability)	GV-1.2
DPA signed with all AI vendors	CC9.2 — Contractual obligations	Art. 26(1) — Obligations of deployers	GV-3.1
No training on customer data	C1.1 — Confidential information protection	Art. 10 — Training, validation, testing data	GV-1.1
Encryption in transit + at rest	CC6.1 — Logical access security	Art. 15 — Accuracy, robustness, cybersecurity	MS-1.1
MFA + SSO enforced for AI tools	CC6.2 — Authentication	Art. 9 — Risk management system	GV-2.1
Sub-processor list reviewed	CC9.2 — Subservice organizations	Art. 28 — Obligations re: third parties	GV-3.1
Breach notification < 72 hours	CC7.3 — Security incident response	Art. 73 — Reporting obligations	MG-1.1
Output accuracy tested	A1.2 — Performance monitoring	Art. 15 — Accuracy requirements	MS-1.1
Bias/fairness assessment (HR/credit)	PI1.5 — Processing integrity	Art. 10(3) — Data governance for high-risk AI	MS-2.1
AI AUP employee training	CC1.4 — Human resources security	Art. 26(6) — Operator training obligations	GV-2.1

EU AI ACT — KEY OBLIGATIONS FOR DEPLOYERS (ART. 26)

EU AI Act: What AI Deployers Must Do

The EU AI Act (effective August 2026 for high-risk AI systems) places obligations on "deployers" — organizations that use AI tools in professional contexts. This summary covers the obligations most relevant to enterprise AI tool adoption.

Art. 26(1)

Use AI in accordance with instructions

Deployers must use AI systems in accordance with the provider's published usage instructions. Using an AI tool for a purpose explicitly excluded in the provider's terms creates liability exposure.

Art. 26(3)

Fundamental rights impact assessment

Deployers of high-risk AI systems (employment, credit, education, law enforcement, critical infrastructure) must conduct a fundamental rights impact assessment before deployment.

Art. 26(5)

Human oversight

Deployers must assign human oversight to high-risk AI systems. Fully automated decisions in high-risk domains without a human review layer are non-compliant under the EU AI Act.

Art. 26(6)

Employee training

Deployers must ensure employees who interact with high-risk AI systems have sufficient AI literacy and training. Documented training is required — informal briefings are insufficient.

Art. 27

Transparency to affected persons

When AI is used to make decisions that significantly affect individuals, those individuals must be informed that a decision was made or assisted by an AI system.

Art. 50

Transparency for general-purpose AI (ChatGPT, Copilot, etc.)

When using general-purpose AI (GPAI) systems, deployers must label synthetic content (AI-generated text, images, audio, video) in contexts where deception risk is high.

APPROVAL DECISION TREE

Approve / Approve-with-Controls / Reject

Use this decision tree after completing the 20-question checklist. Scores are based on questions marked Required (2 points each) and Recommended (1 point each). Maximum score: 30 points.

APPROVE

Score: 27–30 points

All Required questions passed. 2 or fewer Recommended misses.

- All 12 Required questions answered Yes (data exclusion, retention, encryption, MFA, DPA, BAA if applicable, breach notification, sub-processor, IP ownership, incident notification, confidentiality, audit rights)
- Fewer than 3 Recommended questions answered No or Unknown
- Current SOC 2 Type II or ISO 27001 certificate on file
- Documented approval signed by CISO or designated AI Risk Owner

APPROVE WITH CONTROLS

Score: 18–26 points

Some Required misses covered by compensating controls.

- All data exclusion (Q2), encryption (Q6), and breach notification (Q8) questions pass
- Any Required miss must have a documented compensating control (e.g., data types restricted, network isolation, enhanced monitoring)
- Written risk exception signed by CISO and business owner with annual review date
- Tool restricted to non-regulated data classes unless specific BAA/DPA in place
- Quarterly monitoring checkpoint scheduled

REJECT

Score: <18 points or critical fails

Critical Required question failed with no compensating control available.

- Training data exclusion (Q2) not contractually guaranteed — reject outright
- No DPA or AI addendum available — reject until remediated
- Breach notification window exceeds 72 hours with no alternative — reject for regulated data
- Vendor cannot provide any third-party security assessment (SOC 2, ISO, penetration test) — reject for Confidential/Regulated data
- Vendor uses your data for model training by default with no opt-out — reject outright

TOOL-LEVEL QUICK REFERENCE

ChatGPT Enterprise & Microsoft Copilot

ChatGPT (OpenAI) — Enterprise

Training Data Use	Off by default for Enterprise API; configurable in Settings > Data Controls for consumer accounts
Data Retention	Configurable; zero-day retention available for API calls; 30-day default for Enterprise chat
DPA	DPA available at openai.com/policies/data-processing-addendum
BAA (HIPAA)	Available for qualifying Enterprise accounts; request through sales
Encryption	TLS 1.2+ in transit; AES-256 at rest
SOC 2 / ISO	SOC 2 Type II, SOC 3 (public), ISO 27001 certified
Sub-Processors	Published at openai.com/policies/sub-processors (Microsoft Azure, Stripe, others)
EU Data Residency	US-based processing; EU residency not available as of 2025
Key Risk	Consumer accounts: training opt-in by default. Always verify the account tier before approving.

Microsoft Copilot (M365) — Enterprise

Training Data Use	M365 Commercial data is NOT used to train foundation models (Microsoft contractual commitment)
Data Retention	Follows M365 tenant retention policies; admin-configurable
DPA	Covered under Microsoft Online Services Data Protection Addendum (DPA)
BAA (HIPAA)	BAA available through M365 Compliance Center; requires Health Data Services agreement
Encryption	TLS 1.2+ in transit; BitLocker + AES-256 at rest
SOC 2 / ISO	SOC 2 Type II, ISO 27001, ISO 27018 (PII in cloud), FedRAMP Moderate
Sub-Processors	Published at microsoft.com/en-us/download/details.aspx?id=57770
EU Data Residency	EU Data Boundary available for qualifying M365 tenants (EU-based processing)
Key Risk	Copilot inherits M365 permissions — overpermissioned users create data exposure risk via Copilot queries.

TOOL-LEVEL QUICK REFERENCE — CONTINUED

Claude (Anthropic) & Gemini (Google)

Claude (Anthropic) — API / Enterprise

Training Data Use	API: opt-out available (default: Anthropic may use to improve models); Enterprise: contractual exclusion available
Data Retention	API: no persistent storage by default; conversations not stored beyond the session for API calls
DPA	DPA available for Enterprise accounts via sales; Claude.ai consumer DPA available at anthropic.com/legal
BAA (HIPAA)	BAA available for qualifying Enterprise accounts; not available for consumer claude.ai
Encryption	TLS 1.3 in transit; AES-256 at rest
SOC 2 / ISO	SOC 2 Type II (2024); ISO 27001 certification in progress as of 2025
Sub-Processors	Published at anthropic.com/legal/privacy — includes AWS, GCP infrastructure
EU Data Residency	US-based processing; no EU residency option as of 2025
Key Risk	Consumer claude.ai: default training opt-in. API with zero-retention mode is the safest configuration.

Gemini (Google) — Workspace / API

Training Data Use	Workspace: Google contractually commits not to use Workspace data for ads or model training; API: configurable
Data Retention	Workspace: follows Google Workspace admin retention settings; Gemini API: configurable data usage settings
DPA	Cloud Data Processing Addendum (CDPA) at cloud.google.com/terms/data-processing-addendum
BAA (HIPAA)	BAA available for Google Workspace Enterprise Plus and Google Cloud Healthcare API customers
Encryption	TLS 1.3 in transit; AES-256 at rest; key management options via Cloud KMS
SOC 2 / ISO	SOC 2 Type II, SOC 3 (public), ISO 27001, ISO 27018, FedRAMP High (Google Cloud)
Sub-Processors	Published at workspace.google.com/intl/en/terms/subprocessors.html
EU Data Residency	Data regionalization available for Google Workspace Enterprise; EU processing regions supported
Key Risk	Consumer google.com/gemini: data used for model improvement by default. Workspace Enterprise required for enterprise controls.

SHADOW AI — FINDING WHAT YOU DIDN'T APPROVE

Shadow AI Discovery Checklist

85% of employees use at least one AI tool not approved by IT (Gartner 2025). The Samsung incident — where engineers pasted proprietary semiconductor chip source code into ChatGPT — became public because the company had no shadow AI discovery or policy in place. Use this checklist quarterly.

THE SAMSUNG INCIDENT

In April 2023, Samsung engineers used ChatGPT to debug confidential chip manufacturing code — pasting unreleased IP directly into consumer accounts. The data became part of OpenAI's training dataset before the policy was discovered. Samsung banned ChatGPT internally within weeks. 38TB of sensitive data was exposed. Source: Bloomberg, April 2023.

- [ ] Audit browser extensions installed on company devices — AI assistant extensions (ChatGPT, Claude, Gemini, Copilot, Grammarly AI, etc.) are the highest-risk shadow AI vector.
- [ ] Review DNS/proxy logs for connections to: openai.com, api.openai.com, claude.ai, gemini.google.com, bard.google.com, copilot.microsoft.com, chat.mistral.ai, perplexity.ai.
- [ ] Survey department heads quarterly: "Which AI tools is your team using?" Compare responses to the AI inventory — gaps indicate shadow adoption.
- [ ] Check software procurement records for AI-related SaaS subscriptions made with personal or department credit cards (Expensify reports, AP records).
- [ ] Review code repositories (GitHub, GitLab) for hardcoded API keys for AI services — this reveals developer-level shadow usage.
- [ ] Scan outbound email logs for bulk data transfers to AI tool domains — large attachments or forwarding to external AI services is a leading indicator.
- [ ] Monitor Slack/Teams channels for AI tool sharing: employees frequently share tool recommendations in public channels.
- [ ] Check app marketplaces (Microsoft AppSource, Slack App Directory, Salesforce AppExchange) for AI add-ons installed without IT review.

EMPLOYEE AI ACCEPTABLE USE POLICY — TEMPLATE

AI Acceptable Use Policy Template

[COMPANY NAME] — AI Tool Acceptable Use Policy. Version 1.0. Replace all [BRACKET] placeholders with your organization's specifics. This template covers NIST AI RMF GV-1.1 and EU AI Act Art. 26(6) employee training requirements.

1. Purpose

This policy governs the use of artificial intelligence (AI) tools by [COMPANY NAME] employees, contractors, and authorized third parties. It establishes approved tools, data classification requirements, and employee obligations to protect [COMPANY NAME]'s confidential information, customer data, and regulated information.

2. Approved AI Tools

Only AI tools on the [COMPANY NAME] Approved AI Tool List (maintained by [IT/CISO team]) may be used for business purposes. The approved list is available at [link/location]. Employees must not use AI tools not on the approved list for any business purpose. Personal accounts on approved platforms (e.g., consumer ChatGPT free tier) are not approved — only the Enterprise accounts provisioned by IT.

3. Data Classification Restrictions

Employees MUST NOT submit Regulated data (PHI, PII, PCI, ITAR, attorney-client privilege) to any AI tool without explicit written approval from the CISO. Employees MUST NOT submit Confidential data (trade secrets, M&A strategy, unreleased financials, customer lists) to consumer AI tools — only Enterprise-tier approved tools with a signed DPA. Internal data may be submitted to Tier 1 Approved tools only. Public data may be submitted to any approved tool.

4. Employee Obligations

Employees must: (a) verify any AI-generated content before using it in customer-facing communications, legal documents, financial statements, or technical specifications; (b) disclose to recipients when significant portions of a communication or document were AI-generated; (c) report any suspected unauthorized AI use or AI-related security incident to [security@company.com] within 24 hours; (d) complete [COMPANY NAME]'s AI Safety Training (available at [LMS link]) within 30 days of hire and annually thereafter.

AI ACCEPTABLE USE POLICY — CONTINUED

5. Prohibited Uses

Employees must NOT use AI tools for: (a) generating communications that impersonate another person without their knowledge; (b) automated decision-making about another person's employment, credit, benefits, or legal status without human review; (c) creating deepfake audio, video, or imagery of any person without explicit consent; (d) bypassing security controls, generating malware or exploit code, or any other purpose that violates [COMPANY NAME]'s security policies.

6. Vendor Approval Process

Employees who wish to use an AI tool not on the approved list must: (1) submit a Tool Approval Request to [IT/Security team] using the AI Tool Vetting Checklist (available at [link]); (2) wait for written approval before using the tool for any business purpose; (3) not attempt to work around this requirement by using personal accounts or non-company devices for business data.

7. Consequences

Violation of this policy may result in disciplinary action up to and including termination, and may expose [COMPANY NAME] to regulatory penalties, litigation, and reputational harm. Employees who discover a violation must report it immediately — failure to report a known violation is itself a policy violation.

8. Review and Updates

This policy will be reviewed annually by [CISO name/title] and updated to reflect changes in AI tool capabilities, vendor terms, and applicable law. All employees will be notified of material changes within 30 days. Continued employment implies acceptance of the current policy version.

IMPLEMENTATION NOTE

Distribute this policy via your employee handbook platform and require an annual acknowledgment signature. For EU-based employees, ensure the policy references EU AI Act Art. 26(6) obligations. Link the policy to your AI Tool Inventory and Approved Tool List for easy cross-reference.

LIVE TRAINING + IMPLEMENTATION SUPPORT

Put the AI Vetting Kit Into Practice

The kit is the framework. Your team needs hands-on guidance to implement the policy, run the vendor reviews, and train employees before a shadow AI incident turns into a breach.

PERSONAL

$299

1 session · 1 hour

AI tool security briefing: the Samsung case study, 20-question checklist walkthrough, data classification for your specific tools, and shadow AI discovery basics.

Best for: IT managers, security leads, compliance

- 1-hour live AI security briefing
- Samsung case study walkthrough
- Tool vetting checklist review
- Data classification for your stack
- Certificate of completion

EXECUTIVE

$899

1 session · 90 minutes

For CISOs, legal counsel, and AI program leads. Deep dive into NIST AI RMF implementation, DPA negotiation tactics, EU AI Act deployer obligations, and live vendor review for your top 3 AI tools.

Best for: CISOs, General Counsel, DPOs, IT Direct

- 90-minute live session
- NIST AI RMF mapping for your org
- EU AI Act compliance review
- DPA clause negotiation walkthrough
- 3-tool live vetting review
- AI AUP policy customization

BUSINESS

Custom

Unlimited users · 12 months

Full-organization AI security training + quarterly tool vetting reviews. Covers shadow AI discovery, employee AUP rollout, annual vendor reviews, and board briefing materials.

Best for: Organizations rolling out AI governance programs or responding to regulatory scrutiny.

- Unlimited sessions for 12 months
- Shadow AI discovery workshop
- Employee AUP training rollout
- Quarterly vendor review cycle
- Board briefing on AI risk
- EU AI Act deployer compliance review

BOOK YOUR AI SECURITY SESSION

Book at secureveryone.com/book. All sessions include this kit, the AI AUP policy template, and a post-session written summary of your AI tool risk posture.

SecurEveryone LLC · 1201 North Orange Street, Wilmington, DE 19801 · info@secureveryone.com