AI POLICY

Sample Company

Artificial intelligence acceptable use and governance policy

VERSION
1.0

EFFECTIVE
9 October 2026

NEXT REVIEW
9 October 2027

POLICY OWNER
Jordan Reeves, Chief Technology Officer

POSTURE
Balanced

Prepared with the Managed247 AI Policy Generator
ai-policy-generator.managed.co.uk · managed.co.uk

Contents

	Document control	3
1	Purpose and scope	4
2	Our approach to AI	5
3	Definitions	6
4	Roles and responsibilities	8
5	Approved AI tools	10
6	Information you may and may not use with AI	15
7	Permitted uses	17
8	Prohibited uses	18
9	Human oversight, accuracy and accountability	19
10	AI in software development	20
11	Meeting recording and transcription	21
12	AI agents and automation	22
13	AI in recruitment and people decisions	23
14	AI in marketing, communications and content	24
15	Requesting a new AI tool	25
16	Security requirements	26
17	Legal and regulatory compliance	27
18	Training and AI literacy	28
19	Reporting incidents and concerns	29
20	Breaches of this policy	30
21	Review and governance	31
	Appendix A Quick reference for staff	32
	Appendix B AI tool request form	33
	Appendix C Staff acknowledgement	34
	Appendix D Related policies and references	35
	Putting this policy into practice	36

DOCUMENT CONTROL

Sample Company AI Policy

Item	Detail
Document title	Sample Company AI Policy
Version	1.0
Status	Draft for approval by the Board of Directors
Policy owner	Jordan Reeves, Chief Technology Officer
Approved by	The Board of Directors
Effective date	9 October 2026
Review cycle	Every 12 months
Next review	9 October 2027
Applies to	Employees, officers and directors; Contractors, consultants and agency staff; Interns, apprentices and volunteers; Third parties using our systems or acting for us
Classification	Internal
Contact	jordan.reeves@example.com

Version history

Version	Date	Author	Change
1.0	9 October 2026	Jordan Reeves, Chief Technology Officer	First issue

> This policy was produced with the Managed247 AI Policy Generator. It is a template for Sample Company to adopt and adapt, and it is not legal advice. Sample Company should review it with its legal and data protection advisers before adoption, and keep it under review as tools, law and regulatory expectations change.

SECTION 1

Purpose and scope

This policy sets out how Sample Company uses artificial intelligence (AI) tools: which tools are approved, what information may be used with them, how outputs are checked, and the responsibilities of everyone who works with or for us. It exists so that we can take advantage of AI safely, lawfully and with confidence.

AI is already part of everyday software, from email and meeting tools to search and coding assistants. Used well, it saves time, improves quality and opens up new ways of working. Used carelessly, it can leak confidential information, produce confident but wrong answers, infringe other people's rights, or make decisions about people that nobody can explain. This policy draws the line between the two.

It applies alongside our existing policies on information security, data protection, acceptable use of IT and conduct. Where this policy is stricter than another policy, this policy applies to the use of AI.

This policy applies to:

- all employees, officers and directors of Sample Company;
- contractors, consultants, agency and temporary staff working for Sample Company;
- interns, apprentices, volunteers and work-experience placements;
- third parties who use Sample Company systems or information, or who use AI on our behalf, to the extent set out in their contracts.

It covers AI use on company-owned devices and personal devices used for work, and on any system or account used for Sample Company work, including personal subscriptions to AI tools where they are used for work purposes.

It covers all forms of AI, including generative AI (text, code, image, audio and video), AI assistants built into software we already use, AI agents that act on our behalf, automation that uses AI, and AI features within products or services we build, buy or sell.

SECTION 2

Our approach to AI

Sample Company supports the use of AI where it adds value and where the risks are understood and controlled. We maintain a register of approved tools, we are clear about what information can be used with them, and we expect people to apply the same judgement they would apply to any other business tool.

Our default position is that approved tools may be used for everyday work, that conditional tools may be used within their stated limits, and that anything not on the register is not used for Sample Company work until it has been assessed.

Six principles guide every decision in this policy:

Accountable	People are accountable for their work, whatever tools they used to produce it. An AI output is a draft or a suggestion until a person has checked it and chosen to rely on it.
Lawful and fair	We use AI in ways that comply with data protection, equality, consumer, intellectual property and sector rules, and we do not use it to treat people unfairly.
Secure	We protect our information, our customers' information and our systems. We only use tools that meet our security and contractual requirements, and we never put credentials or restricted data into unapproved tools.
Transparent	We are honest about where AI has been used when it matters to the person receiving the output, and we can explain how an AI-assisted decision was reached.
Human-led	AI supports people. It does not replace human judgement in decisions that affect customers, colleagues or the business.
Proportionate	Controls are proportionate to risk. Low-risk, everyday uses are made easy; higher-risk uses attract more scrutiny.

SECTION 3

Definitions

Artificial intelligence (AI)	Software that performs tasks which would normally need human intelligence, such as generating text, images or code, summarising information, recognising patterns, making predictions or taking actions.
Generative AI	AI that creates new content (text, code, images, audio or video) in response to a prompt. Examples include ChatGPT, Claude, Microsoft 365 Copilot and Gemini.
AI assistant	A generative AI tool used interactively by a person, including assistants built into software such as Microsoft 365, Google Workspace, GitHub or a CRM.
AI agent	An AI system that can plan and carry out multi-step tasks with limited supervision, such as browsing websites, filling in forms, sending messages, running code or acting within business systems on a person's behalf.
Approved tool	An AI tool listed in the AI tools register in this policy with a status of Approved or Conditional, used in the version and under the conditions stated there.
Company version	A version of an AI tool that Sample Company provides and pays for, on a business, team or enterprise plan, signed in with a work email (for example ChatGPT Business, Claude Team or Microsoft 365 Copilot). The provider's contract commits that our data is not used to train its models, is handled under appropriate data protection terms and can be administered by Sample Company. A free or personal login to the same tool is not the company version.
Free or personal version	A free, trial or personally paid version of an AI tool, signed in with a login an individual set up for themselves. Data entered into these versions may be retained, reviewed or used for training by the provider, and Sample Company has no contractual control over it.
Prompt	Anything entered into or uploaded to an AI tool, including text, files, images, audio, screenshots and data connections.
Output	Anything an AI tool produces, including text, code, images, summaries, decisions, actions and recommendations.
Personal data	Information relating to an identified or identifiable living person, as defined by UK GDPR and the Data Protection Act 2018.

Confidential information	Information that Sample Company or a third party would not want disclosed, including commercial, financial, technical, legal and client information, whether or not it is marked confidential.
Shadow AI	AI tools, or free and personal versions of tools, used for work without approval under this policy, including personal subscriptions.

SECTION 4

Roles and responsibilities

Everyone has a part to play. The specific responsibilities are:

The Chief Technology Officer (Jordan Reeves)

- Owns this policy, keeps the AI tools register current and approves or declines requests to use new tools and new uses.
- Reports to the Board of Directors on AI use, incidents and risks, and chairs the AI working group where one exists.
- Decides, with the data protection lead, when a data protection impact assessment is needed.

IT and information security

- Configure approved tools (single sign-on, enterprise tenancy, data controls, logging) and remove access when people leave or change role.
- Monitor for unapproved AI use and maintain technical controls such as data loss prevention and web filtering.
- Carry out security reviews of new tools and respond to AI-related incidents.

Line managers

- Make sure their teams know this policy and have completed the required training.
- Apply the human review checkpoints to work produced in their area and set a good example in their own use of AI.
- Raise requests for new tools and report incidents promptly.

Everyone covered by this policy

- Use only approved tools, within their stated conditions, and only with information this policy allows.
- Check AI outputs before relying on them and remain responsible for the result.
- Complete AI training when asked, report incidents and concerns, and ask before trying anything this policy does not cover.

People team (HR)

- Include this policy in onboarding and keep training records.
- Advise on the use of AI in recruitment, performance and employee relations, and handle breaches under the disciplinary procedure.

Data protection lead

- Advise on lawful basis, data protection impact assessments, individual rights and transparency where personal data is involved.

- Maintain the record of processing activities for AI uses that involve personal data.

Procurement and legal

- Make sure contracts with AI providers, and with suppliers who use AI on our behalf, contain the required terms.
- Review intellectual property, confidentiality and liability questions raised by new uses.

SECTION 5

Approved AI tools

The register below lists the AI tools Sample Company has assessed. A tool that is not listed is not approved for Sample Company work until it has been through the request process in this policy. The register is maintained by the Chief Technology Officer (Jordan Reeves); the version in this policy is current as at 9 October 2026.

Each status means:

- Approved: may be used for everyday Sample Company work by the functions listed, in the version shown, within the data rules in this policy.
- Conditional: may be used only under the conditions shown, by the functions listed, and in the version shown. If you cannot meet the conditions, do not use the tool.
- Not approved: must not be used for Sample Company work or with Sample Company information, on any device or in any version.
- Under review: not yet approved. Do not use it for Sample Company work until the status changes. Contact the Chief Technology Officer (Jordan Reeves) if you need it.

Tool	Status	Version	Who may use it	Conditions
General assistants				
ChatGPT (OpenAI)	Conditional	Company version	All staff	ChatGPT Business or Enterprise workspace only; training on our data switched off; nothing classed as never.
Claude (Anthropic)	Conditional	Company version	All staff	Claude Team or Enterprise plan only; nothing classed as never.
Microsoft 365 Copilot (Microsoft)	Approved	Company version	All staff	Within the company Microsoft 365 tenant only; it respects existing document permissions, so keep those tidy.
Microsoft Copilot (consumer and free Copilot Chat) (Microsoft)	Not approved	None	No one	Do not use for company work.
Google Gemini (Google)	Conditional	Company version	All staff	Google Workspace business licence only; personal Gmail accounts are not permitted.
Perplexity (Perplexity AI)	Conditional	Company version	All staff	Enterprise Pro only; research and public information; verify sources before use.
Grok (xAI)	Not approved	None	No one	Do not use for company work.
Meta AI (Meta)	Not approved	None	No one	Do not use for company work.
DeepSeek	Not approved	None	No one	Do not use for company work.
Le Chat (Mistral AI)	Under review	None	No one	Do not use until approved.

Tool	Status	Version	Who may use it	Conditions
Sample Assistant (In-house)	Conditional	Any version	All staff	Internal pilot only.
Coding assistants				
GitHub Copilot (GitHub (Microsoft))	Approved	Company version	IT and Developers	Copilot Business or Enterprise with the public code filter on and no training on our code.
Claude Code (Anthropic)	Conditional	Company version	IT and Developers	Team or Enterprise plan; least-privilege repositories; no production credentials in its environment.
Cursor (Anysphere)	Conditional	Company version	Developers	Business plan with privacy mode on.
OpenAI Codex (OpenAI)	Conditional	Company version	Developers	Business or Enterprise workspace only; sandboxed repositories.
Amazon Q Developer (Amazon Web Services)	Conditional	Company version	IT and Developers	Pro tier within the company AWS organisation.
Gemini Code Assist (Google)	Conditional	Company version	Developers	Enterprise edition only.
Tabnine	Under review	None	No one	Do not use until approved.
Replit	Under review	None	No one	Do not use until approved.
Lovable	Under review	None	No one	Do not use until approved.
Bolt (StackBlitz)	Under review	None	No one	Do not use until approved.
Meeting assistants and notetakers				
Microsoft Teams meeting recap (Copilot) (Microsoft)	Approved	Company version	All staff	Company tenant only; participants informed; retention per the meeting rules.
Zoom AI Companion (Zoom)	Conditional	Company version	All staff	Company Zoom account only; participants informed.
Google Meet notes with Gemini (Google)	Conditional	Company version	All staff	Company Workspace only; participants informed.
Otter.ai	Not approved	None	No one	Do not use for company work.
Fireflies.ai	Not approved	None	No one	Do not use for company work.
Fathom	Not approved	None	No one	Do not use for company work.
Read AI	Not approved	None	No one	Do not use for company work.
tl;dv	Not approved	None	No one	Do not use for company work.
Granola	Not approved	None	No one	Do not use for company work.
Image, video and audio generation				
Adobe Firefly (Adobe)	Approved	Company version	Marketing	Creative Cloud for Teams; commercially safe for published work.

Tool	Status	Version	Who may use it	Conditions
Canva AI (Canva)	Conditional	Company version	Marketing and Sales	Canva Teams; brand review before publication.
Midjourney	Conditional	Company version	Marketing	Pro or Mega plan with stealth mode; internal concepts only unless IP review.
ChatGPT image generation (OpenAI)	Conditional	Company version	All staff	Within the approved ChatGPT workspace; no real people.
Sora (OpenAI)	Under review	None	No one	Do not use until approved.
Runway	Under review	None	No one	Do not use until approved.
ElevenLabs	Conditional	Company version	Marketing and Sales	No voice cloning without written consent; synthetic voices disclosed.
Synthesia	Conditional	Company version	Marketing and Sales	Stock avatars or consented likeness only; labelled as AI-generated.
HeyGen	Under review	None	No one	Do not use until approved.
Stable Diffusion (Stability AI)	Under review	None	No one	Do not use until approved.
Writing, productivity and translation				
Grammarly	Conditional	Company version	All staff	Business plan; the extension is disabled on confidential systems.
Notion AI (Notion)	Conditional	Company version	All staff	Within the company Notion workspace only.
Jasper	Under review	None	No one	Do not use until approved.
Copy.ai	Under review	None	No one	Do not use until approved.
Gamma	Conditional	Company version	All staff	Team plan; public and internal information only.
Microsoft Designer (Microsoft)	Conditional	Company version	All staff	Signed in with the work account.
Adobe Acrobat AI Assistant (Adobe)	Conditional	Company version	All staff	Enterprise licence; no documents classed as never.
DeepL	Conditional	Company version	All staff	DeepL Pro only; the free version is not permitted for company text.
Google Translate (Google)	Conditional	Free version allowed	All staff	Public information only.
Agents, browser agents and automation				
ChatGPT agent mode (OpenAI)	Not approved	None	No one	Do not use for company work.
Claude in Chrome and computer use (Anthropic)	Not approved	None	No one	Do not use for company work.
Perplexity Comet (Perplexity AI)	Not approved	None	No one	Do not use for company work.
Manus	Not approved	None	No one	Do not use for company work.

Tool	Status	Version	Who may use it	Conditions
Microsoft Copilot Studio agents (Microsoft)	Conditional	Company version	IT	Built and published by IT only; least privilege; logged.
Zapier and Make AI automations (Zapier and Make)	Conditional	Company version	Operations	IT-managed logins only; no AI steps on personal or confidential data.
n8n AI workflows (n8n)	Under review	None	No one	Do not use until approved.
In-house agents built on model APIs (Various)	Conditional	Any version	IT and Developers	Risk assessment and sign-off from the AI lead before any production use.
AI built into business platforms				
Salesforce Agentforce and Einstein (Salesforce)	Conditional	Company version	Marketing and Sales	Enable features only after IT review of data access.
HubSpot Breeze (HubSpot)	Conditional	Company version	Marketing and Sales	Enable features only after IT review.
ServiceNow Now Assist (ServiceNow)	Conditional	Company version	IT	Enable features only after IT review.
Zendesk AI (Zendesk)	Conditional	Company version	Customer service	The customer-facing AI rules apply.
Atlassian Rovo (Atlassian)	Conditional	Company version	All staff	Enable only after IT review of data access.
Slack AI (Salesforce)	Conditional	Company version	All staff	Company workspace only.
Workday AI (Workday)	Conditional	Company version	HR and Finance	The people-decision rules apply; HR and data protection approval.
Sage Copilot (Sage)	Conditional	Company version	Finance	Finance reviews every output before use.
Xero JAX (Xero)	Conditional	Company version	Finance	Finance reviews every output before use.
Intercom Fin (Intercom)	Conditional	Company version	Customer service	The customer-facing AI rules apply.

Versions of a tool

- Company version: only the version Sample Company provides and pays for, signed in with a work login. Free and personal versions must not be used.
- Free version allowed: the free or a personal version is acceptable, subject to the data rules.
- Any version: any version of the tool, subject to the data rules.

Where a tool offers a setting to prevent our data being used to train the provider's models, it must be switched on. Where a tool allows chat history or uploaded files to be retained, retention must be set to the shortest practical period.

Approved tools are used in the company version, signed in with a Sample Company work login protected by single sign-on and multi-factor authentication wherever the tool supports it. Free and

personal versions, even of the same tool, are not approved.

AI features that appear inside software we already use (for example in a CRM, finance system or productivity suite) are covered by this register. If a supplier switches on an AI feature that is not listed, do not use it until it has been assessed.

SECTION 6

Information you may and may not use with AI

The single biggest risk with AI tools is information leaving Sample Company without control. The rules below say what may be used with which tools. They apply to prompts, uploaded files, screenshots, voice input, meeting recordings and any data connection you give a tool.

Information class	Examples	Rule
Public information	Published website content, press releases, marketing material already in the public domain.	Any approved tool
Internal business information	Internal guidance, process documents, non-sensitive meeting notes, general business correspondence.	Company version only
Confidential and commercially sensitive	Strategy, pricing, contracts, unreleased plans, supplier terms, board papers, anything marked confidential.	Company version only
Personal data	Names, contact details, HR records, customer records, anything that identifies a living person.	Company version only
Special category and highly sensitive data	Health, ethnicity, religion, sexuality, politics, trade union membership, biometrics, criminal records, children's data.	Never
Client and customer confidential information	Client documents, data processed on behalf of clients, anything received under a non-disclosure agreement.	Company version only
Source code and technical designs	Source code, architecture diagrams, infrastructure configuration, API designs.	Company version only
Financial records and price-sensitive information	Management accounts, forecasts, payroll, unreleased results, price-sensitive information.	Company version only
Credentials and security configuration	Passwords, API keys, tokens, certificates, security settings, network diagrams showing controls.	Never

What each rule means

- Any approved tool: May be used with any approved tool, in the version the register shows for that tool and within its conditions.
- Company version only: Approved tools used in the company version only, never the free or personal version, because the company contract prevents the provider training on our data. Use the minimum needed for the task.
- Never: Not to be entered into any AI tool, under any licence, unless the Chief Technology Officer (Jordan Reeves) and the data protection lead have approved a specific, documented exception.

Before you enter information, remove or replace what the tool does not need. A redacted extract is usually enough. If you are unsure which class information falls into, treat it as the higher class.

Information generated by an AI tool from Sample Company data inherits the classification of the data that went in.

Where a tool stores or processes Sample Company data, we prefer providers that keep it in the United Kingdom, and we require that it is not transferred outside the UK without appropriate safeguards.

SECTION 7

Permitted uses

The following uses of approved tools are permitted and encouraged, subject to the data rules and the human review checkpoints:

- Drafting, editing and improving documents, emails, presentations and other written material.
- Summarising documents, reports, threads and meetings so that people can get to the point faster.
- Research, explanation and learning, provided facts and sources are checked before they are relied upon.
- Writing, explaining, reviewing, refactoring and testing code, within the software development rules in this policy.
- Analysing and visualising data that the data rules allow, and producing first-pass insights for a person to validate.
- Drafting responses to customers and suppliers for a person to review before sending.
- Generating ideas, copy, images and campaign material for marketing, subject to brand and legal review.
- Translating and localising content, with a fluent reviewer for anything customer-facing or contractual.
- Brainstorming, planning, structuring problems and preparing for meetings.
- Automating repetitive administrative tasks such as formatting, categorising, data entry from approved sources and template filling.
- Producing meeting notes and actions with an approved meeting assistant, under the meeting rules in this policy.
- Personal upskilling, practising skills and exploring how AI could improve our work, using approved tools and non-sensitive information.

SECTION 8

Prohibited uses

The following are prohibited for everyone covered by this policy:

- Using AI to make or materially influence a decision with legal or similarly significant effects on a person (for example recruitment screening, dismissal, performance ratings, credit, or pricing for an individual) without meaningful human review by someone with the authority to overturn the result.
- Using any AI tool that is Not approved or Under review for Sample Company work, or entering Sample Company information into it.
- Using the free, personal or consumer version of an AI tool for Sample Company work where this policy requires the company version.
- Entering passwords, keys, tokens, security configuration or any information classed as never in the data rules into an AI tool.
- Creating or sharing synthetic images, video or audio of real people (including colleagues, customers or public figures) without their documented consent, or any content intended to deceive.
- Using AI to create, test or distribute malicious code, phishing content, or anything designed to defeat security controls, other than authorised security testing.
- Presenting AI-generated legal, financial, medical, tax or other professional advice to a client or customer without review by an appropriately qualified person.
- Uploading client documents, information received under a non-disclosure agreement, or third-party confidential information to any tool unless the data rules and the relevant contract allow it.
- Using AI to monitor, profile or assess colleagues covertly, or in any way that has not been approved by HR and the data protection lead.
- Generating, using or circulating content that is discriminatory, harassing, defamatory, obscene or otherwise contrary to our conduct policies.
- Circumventing technical controls, logging or approval processes that Sample Company has put in place around AI tools.
- Presenting unverified AI output as checked fact, as original work where originality matters, or as the position of Sample Company without authorisation.
- Using AI tools to collect personal data or content from websites, social media or other sources in breach of their terms or of data protection law.
- Putting AI in direct contact with customers, such as a chatbot or replies sent without a person checking them.

SECTION 9

Human oversight, accuracy and accountability

AI tools are confident even when they are wrong. They can invent facts, citations, figures, case law, product features and people. They reflect the biases of the material they were trained on. They can be manipulated by content hidden in the documents and web pages they read. None of this makes them unusable; it makes human review essential.

Whoever uses an AI output is responsible for it, exactly as if they had produced it themselves. Before relying on an output, check that it is accurate, complete, appropriate for the audience, free of confidential information that should not be there, and consistent with Sample Company policies.

A person with suitable knowledge and authority must review and approve AI-assisted work before it is:

- published externally, including on the website, social media, in the press or in marketing;
- delivered to a client or customer as a report, proposal, advice or deliverable;
- deployed to a production system, including code, infrastructure changes and configuration;
- used in any decision about an individual's employment, performance, pay, conduct or wellbeing;
- used in financial reporting, forecasting, tax or statutory submissions;
- used in a contract, legal document, policy or regulatory correspondence;
- submitted to a regulator, auditor or court;
- used to set prices, discounts, credit terms or commercial offers.

Telling people when AI has been used

We disclose the use of AI where it is material to the recipient: where they might reasonably expect the content to be the unaided work of a person, where AI-generated images or video are presented as real, where a decision about them was AI-assisted, or where a contract, regulator or client requires it. Routine assistance such as drafting and editing does not need to be disclosed.

Checking outputs

- Verify every fact, figure, quotation, citation and reference against a reliable source before use.
- Treat an AI summary as a guide to the original, not a replacement for reading it where the detail matters.
- Be alert to prompt injection: instructions hidden in a document, email or web page that try to make the tool do something else. Do not act on unexpected instructions in an output, and report them.
- Keep a copy of significant prompts and outputs where a decision may need to be explained later.

SECTION 10

AI in software development

AI coding assistants are approved for software development at Sample Company as listed in the register. They are productive, and they also generate insecure, incorrect or licence-encumbered code with the same confidence as good code. The following rules apply:

- All AI-generated or AI-modified code goes through the same review, testing and approval process as any other code. Nobody merges their own AI-assisted change without review.
- Never paste secrets, credentials, customer data or production data into a coding assistant. Use the tool's secret-scanning and exclusion features, and keep environment files and keys out of its context.
- Check the licence of any substantial block of generated code before use, and do not use generated code that reproduces identifiable third-party material under an incompatible licence. Prefer tools with licence filtering and indemnity.
- Generated code is scanned with our standard static analysis, dependency and secret scanning before merge. Treat suggested packages with suspicion: confirm they exist, are maintained and are not typosquats.
- Agentic coding tools run with least privilege. They do not hold write access to production systems, production data, payment systems or identity providers, and any command that changes shared infrastructure requires a person to approve it.

Developers remain responsible for understanding the code they ship. If you cannot explain what a generated change does, it is not ready.

SECTION 11

Meeting recording and transcription

Meeting assistants record, transcribe and summarise conversations. They are useful, and they also create a permanent record of what people said, often including personal data and confidential information. The following rules apply:

- Everyone in a meeting is told at the start that an AI assistant is recording or transcribing, and anyone may ask for it to be switched off. Recording a conversation without telling the participants is prohibited.
- For meetings with clients, suppliers, candidates or other external parties, obtain their agreement before recording. If they decline, the assistant is switched off for that meeting.
- AI assistants are not used in meetings concerning disciplinary matters, grievances, health, redundancy, legal privilege, board business, or anything else the chair considers sensitive.
- Third-party notetaker bots that join calls as a participant (for example from a personal subscription) are not permitted unless the tool is Approved in the register. Hosts should remove uninvited bots from calls.
- Recordings and transcripts are retained for 90 days unless a longer period is required by our records policy or a legal hold, after which they are deleted. They are stored only in approved, access-controlled locations.

Summaries produced by an assistant are checked by the meeting owner before being circulated, and are not treated as formal minutes unless approved as such.

SECTION 12

AI agents and automation

AI agents can browse the web, operate software, send messages, make purchases and change data with limited supervision. They multiply both the benefit and the risk of AI, because a mistake or a manipulated instruction can be acted on immediately.

AI agents may be used only in approved pilots sponsored by the Chief Technology Officer (Jordan Reeves), with a defined scope, a named owner, documented controls and a review date. Outside an approved pilot, agent features in approved tools remain switched off.

The following controls apply to any agent that is permitted:

- A person approves any irreversible or outward-facing action before it happens: sending external communications, making payments or purchases, deleting or overwriting data, changing access or permissions, or publishing content.
- Agents run under their own identity with the minimum access needed, never under a person's full account, and never with standing administrative rights.
- Every action an agent takes is logged with the prompt or trigger that caused it, and the logs are retained and reviewable.
- Every agent can be stopped immediately by its owner and by IT, and is stopped by default if it behaves unexpectedly.
- Agents do not have write access to production systems, financial systems, identity providers or customer data without a specific risk assessment approved by the Chief Technology Officer (Jordan Reeves).

Connecting an AI tool to a Sample Company system (for example through a connector, plug-in, API key or integration platform) is treated as deploying an agent and follows the request process in this policy.

SECTION 13

AI in recruitment and people decisions

AI is increasingly offered within recruitment, HR and workforce tools. Decisions about people carry legal, ethical and reputational risk, and the following rules apply:

- AI does not make or materially determine decisions about hiring, promotion, pay, performance, conduct or dismissal. It may assist with administration, drafting and summarising, and a qualified person makes every decision.
- Candidates and employees are told where AI is used in a process that affects them, what it does, and how to request human review.
- Any AI used in people processes is assessed for bias and adverse impact before use and at least annually, with the results recorded.
- Outputs that describe or rank individuals (for example CV screening, sentiment or productivity analysis) are reviewed by HR before being acted on and are retained only as long as necessary.

Covert monitoring of staff through AI is prohibited. Any workforce analytics using AI require approval from HR and the data protection lead, and consultation where required.

SECTION 14

AI in marketing, communications and content

Generative AI is a powerful creative tool and its outputs are public-facing by nature. The following rules apply to marketing, communications and content:

- AI-generated copy and imagery go through the normal brand and editorial review before publication, and must be accurate, substantiated and compliant with advertising rules.
- AI-generated or materially AI-edited images, video and audio are labelled where a reasonable person could mistake them for real events, products or people, and are never used to misrepresent our products or results.
- AI is never used to create fake reviews, testimonials, case studies, endorsements, social media accounts or engagement.
- The likeness, voice or style of a real person is not generated without their written consent, and we do not imitate the distinctive work of identifiable artists or brands.
- Outputs are checked for third-party trademarks, logos, copyrighted material and anything that could infringe others' rights before use, and we prefer tools with commercial-use terms and indemnity for published material.

The marketing lead is accountable for AI-assisted content published in the name of Sample Company.

SECTION 15

Requesting a new AI tool

Anyone can request a new AI tool, a new use of an existing tool, or a change to the register. Requests go to the IT service desk and are assessed by the Chief Technology Officer (Jordan Reeves) with IT and, where relevant, the data protection lead, procurement and legal. The steps are:

1. Submit a request describing the tool, what it will be used for, who will use it, what information it will handle and why existing approved tools are not suitable.
2. IT reviews the tool's security: authentication and single sign-on, access control, data handling, encryption, logging, admin controls, and the provider's security certifications.
3. Where personal data is involved, the data protection lead decides whether a data protection impact assessment is required and records the lawful basis and any transfer safeguards.
4. Procurement and legal confirm the contract terms: no training on our data, data residency, retention and deletion, sub-processors, confidentiality, intellectual property in outputs, liability and exit.
5. the Chief Technology Officer (Jordan Reeves) decides the status and conditions, updates the register, and IT configures the tool. The requester is told the outcome and any conditions.

A tool may be approved for a pilot with a limited group before a wider decision. Tools are re-assessed when the provider changes its terms, features or ownership, when an incident occurs, and at least annually.

Free trials and personal subscriptions are not a shortcut: the same assessment applies before Sample Company information is used.

SECTION 16

Security requirements

The following security requirements apply to all AI use at Sample Company:

- Use approved tools through your Sample Company work login, with single sign-on and multi-factor authentication. Do not share logins or API keys.
- Keep AI tools and browser extensions up to date and install them only from approved sources. IT may block unapproved AI tools and extensions at the network, device and browser level.
- Do not grant an AI tool access to your mailbox, files, calendar, code repositories or business systems unless the connection has been approved. Review and remove connections you no longer need.
- Treat AI output as untrusted input: do not run generated commands, scripts or code on company systems without understanding them, and never disable security tools because an AI suggests it.
- Be alert to prompt injection and social engineering through AI: content in emails, documents and websites can attempt to hijack a tool you are using. Report anything unexpected.
- Report suspected data exposure through an AI tool immediately, including accidental pasting of restricted data, so that it can be contained.

Monitoring and technical controls

Sample Company operates the following controls, in line with our monitoring and privacy notices:

- Data loss prevention controls monitor, and may block, the transfer of restricted information to AI tools.
- IT monitors for unapproved AI tools and services (shadow AI) and will contact people using them to move to approved alternatives.
- Use of the company versions of approved tools is logged and may be reviewed for security, compliance and incident investigation.
- Access to AI tools is reviewed as part of our regular access reviews and removed promptly when people leave or change role.

SECTION 17

Legal and regulatory compliance

The use of AI by Sample Company must comply with the law and with our contractual obligations. The main areas are:

Data protection

UK GDPR and the Data Protection Act 2018 apply whenever personal data is used with or produced by AI. We need a lawful basis, we must be transparent with the people concerned, we carry out a data protection impact assessment for higher-risk uses, we honour individual rights including the rights around automated decision-making, and we apply appropriate safeguards to any international transfer. Information Commissioner's Office guidance on AI is followed.

Equality

The Equality Act 2010 prohibits discrimination, including discrimination that arises from a biased algorithm. We test and monitor AI used in decisions about people for adverse impact.

Intellectual property

Copyright in AI outputs is uncertain in law. We check outputs for infringement of third-party rights, we are careful about relying on AI outputs for material where ownership matters (for example brand assets and product code), and we use provider terms that give us commercial-use rights and, where possible, indemnity. Confidential information and trade secrets lose protection if disclosed to an uncontrolled tool.

Confidentiality and contracts

Client contracts, non-disclosure agreements and supplier terms may restrict or prohibit the use of AI with their information or in delivering services to them. Check the contract before using AI on client work, and tell clients where our contracts require it.

Consumer protection and advertising

Marketing claims generated with AI must be accurate and substantiated, and comply with the CAP Code, consumer protection law and the Digital Markets, Competition and Consumers Act 2024, including the rules on fake reviews.

Records

Significant prompts, outputs and AI-assisted decisions are retained in line with our records management policy so that decisions can be explained and audited.

SECTION 18

Training and AI literacy

Safe use of AI depends on people understanding what the tools can and cannot do.

- Everyone covered by this policy completes AI awareness training on joining, covering this policy, the data rules, the approved tools and how to spot and report problems.
- Refresher training is completed at least annually and whenever this policy changes materially.
- Role-specific training is provided for people who use AI in higher-risk areas, including software development, HR and recruitment, marketing, customer-facing services and finance.
- Training supports the AI literacy expectations of the EU AI Act where it applies, and records of completion are kept by HR.

Managers make time for training, and the Chief Technology Officer (Jordan Reeves) keeps the material current as tools and risks change.

SECTION 19

Reporting incidents and concerns

An AI incident is any event where AI use has, or could have, caused harm to Sample Company, our customers, our people or a third party. Examples include:

- Restricted, personal or confidential information entered into an unapproved tool, or into an approved tool outside the data rules.
- An AI output containing errors, fabricated content or confidential information that was published, sent or acted upon.
- An AI agent or automation taking an unintended action, such as sending, deleting or purchasing something.
- A suspected prompt injection, account compromise or malicious use of an AI tool.
- Discovery of an unapproved AI tool in use with Sample Company information, or a supplier using AI on our data without agreement.
- A customer, candidate or employee complaint about an AI-assisted decision or interaction.

Report incidents within 24 hours to the IT service desk. Do not try to delete evidence or quietly fix the problem first: speed matters more than tidiness, and reporting a genuine mistake promptly is always the right thing to do.

Incidents involving personal data are also handled under our data breach procedure, which includes assessing whether notification to the Information Commissioner's Office is required within 72 hours and whether affected individuals must be told.

IT and the Chief Technology Officer (Jordan Reeves) contain the incident, record it, and decide what needs to change: the register, the controls, the training or this policy.

Concerns about how AI is being used, including by colleagues, suppliers or Sample Company itself, can be raised with the Chief Technology Officer (Jordan Reeves), with HR, or through the whistleblowing policy. Concerns raised in good faith are always welcome.

SECTION 20

Breaches of this policy

Breaches of this policy are taken seriously. Depending on the circumstances they may be dealt with under the disciplinary procedure, up to and including dismissal, and for contractors and third parties may result in termination of the engagement. Deliberate or reckless disclosure of restricted information, and attempts to circumvent controls, are treated as gross misconduct.

We recognise that AI is new to many people. Honest mistakes that are reported promptly are treated as learning opportunities, and nobody will be penalised for raising a concern in good faith.

SECTION 21

Review and governance

the Chief Technology Officer (Jordan Reeves) owns this policy and reports to the Board of Directors on AI use at least quarterly. An AI working group, with representatives from IT, information security, HR, legal and the main business functions, meets regularly to review requests, incidents, metrics and emerging risks.

This policy is reviewed every 12 months, and sooner if there is a significant change in the tools we use, the law, our regulators' expectations, or following an incident. The next scheduled review is 9 October 2027.

We track:

- the number and status of tools on the register and open requests;
- training completion;
- incidents reported, their causes and time to resolve;
- adoption of approved tools and retirement of shadow AI;
- outcomes of bias and accuracy monitoring for higher-risk uses.

Changes to this policy are recorded in the version history and communicated to everyone covered by it.

APPENDIX A

Quick reference for staff

Keep this page to hand. It is a summary; the policy itself is the rule.

Do	Do not
• Use the approved tools in the register, in the company version, signed in with your work login.	• Do not use the free or personal version of an AI tool for work where the company version is required.
• Check what information you are allowed to use with each tool, and remove what the tool does not need.	• Do not enter passwords, keys, special category personal data or anything marked never in the data rules into any AI tool.
• Check every output before you rely on it: facts, figures, sources, tone and confidentiality.	• Do not paste client documents or non-disclosure material into a tool unless the data rules and the contract allow it.
• Get the right person to review AI-assisted work before it goes to a customer, into production or into a decision about a person.	• Do not let AI make decisions about people, prices or money without a person who can overturn it.
• Tell people when an AI assistant is in a meeting, and when AI has materially shaped something they receive, where this policy requires it.	• Do not create deepfakes, fake reviews or content that imitates real people.
• Report mistakes and anything odd straight away to the IT service desk.	• Do not run generated code or commands you do not understand on company systems.
• Ask the Chief Technology Officer (Jordan Reeves) if you want to use a new tool or try a new use.	• Do not connect AI tools to Sample Company systems or data without approval.

APPENDIX B

AI tool request form

Use this form, or the equivalent in our service desk, to request a new AI tool or a new use of an existing tool.

Tool name and provider

Requested by, team and date

What it will be used for

Who will use it

Information it will handle (data classes from this policy)

Version or plan, and cost

Why existing approved tools are not suitable

Links to the provider's terms, security and data protection information

Decision, status and conditions, completed by the Chief Technology Officer (Jordan Reeves)

Review date

APPENDIX C

Staff acknowledgement

I confirm that I have read and understood the Sample Company AI Policy (version 1.0, effective 9 October 2026), that I will use AI tools only as it allows, and that I will report incidents and concerns promptly.

Name _____

Role _____

Signature _____

Date _____

APPENDIX D

Related policies and references

This policy should be read with:

- IT Acceptable Use Policy
- Data Protection Policy and privacy notices
- Information Security Policy
- Data Breach Procedure
- Disciplinary Procedure and Code of Conduct
- Procurement and Supplier Management Policy
- Records Management and Retention Policy

External references:

- Information Commissioner's Office: guidance on AI and data protection
- National Cyber Security Centre: guidance on AI and secure development
- UK Government: AI Playbook and Generative AI Framework for HM Government
- EU AI Act: Regulation (EU) 2024/1689
- Competition and Markets Authority: guidance on fake reviews and consumer law

MANAGED247

Putting this policy into practice

A policy is only as good as the controls behind it. Managed247 is a UK managed IT and cyber security provider, and the items below are the ones we are most often asked to help with after an AI policy is adopted.

- Microsoft 365 Copilot readiness: permissions clean-up, sensitivity labels and a measured rollout so Copilot only surfaces what each person should see.
- Data loss prevention and information protection: policies that stop restricted information reaching unapproved AI tools, with the evidence to show it.
- Shadow AI discovery: finding the AI tools already in use across the estate and moving people onto approved alternatives.
- Identity and access: single sign-on, multi-factor authentication and conditional access for every approved AI tool.
- Security awareness and AI literacy training for staff, with records kept for audit.
- Managed detection and response, so an AI-related incident is spotted and contained quickly.

To talk any of this through, email hello@managed.co.uk or visit managed.co.uk. Replying to the email that delivered this policy comes straight to us.