inology I.T.

AI POLICY PACK

Use AI like a sharp tool, not a lazy shortcut.

Five pieces. One bundle. Everything your team needs to put AI to work without putting your business at risk.

Prepared by · Inology IT
Version 1.0 · May 2026
inology.co.uk · 0161 503 3536

How to use this pack

AI is no longer optional. Used well, it sharpens every team in your business. Used carelessly, it leaks data, generates fiction as fact, and creates legal exposure that nobody planned for. This pack is the bridge.

Inside you will find five working documents. Each can be lifted out and used on its own:

1. AI Acceptable Use Policy — the master 27-section policy for staff to read and sign.
2. Data Classification Matrix — a single-page reference: what's safe, what needs anonymising, what's off-limits.
3. Approved Tools Register — a template you fill in, vetting AI tools before staff use them.
4. Incident Response Checklist — what to do in the first 30 minutes after a leak.
5. Staff One-Pager — print-and-pin reminder for every desk.

How to roll it out

- Read this pack end-to-end before circulating it.
- Customise the Approved Tools Register and the Incident Response contact details to match your environment.
- Brief the team in a 30-minute session — the One-Pager and the Data Classification Matrix are designed for that meeting.
- Get every member of staff to sign the acceptance page in Part 1.
- Review every six months. AI moves quickly; this policy is a living document.

Questions? Get in touch — info@inology.co.uk · 0161 503 3536.

PART 1

AI Acceptable Use Policy

The master document. Twenty-seven sections covering scope, classification, tools, incidents, ethics and acceptance. Read it. Sign it. Live by it.

inology I.T.

Artificial Intelligence Acceptable Use Policy

Version	1.0
Effective Date	01 May 2026
Document Owner	Inology IT
Review Cycle	Every 6 months

1. Our approach to AI

AI offers real opportunities — leverage, creativity, efficiency. We want our team and our clients' teams to use these tools to do their best work.

But these tools also introduce new risks: data leakage, copyright exposure, hallucinated facts presented as truth. Risks that, mishandled, can damage a business and its reputation in days, not months.

This policy is the roadmap for how we work with AI safely. The goal is simple: protect the business over the long term so it can continue to innovate without putting hard-won trust at risk.

2. Personal responsibility

Every member of staff is responsible for understanding and applying this policy.

If something is unclear, ask before you act. "I didn't know" is not a defence when client data is involved.

3. Key definitions

- Generative AI — Artificial intelligence capable of creating new content (text, images, code) in response to prompts. Examples: ChatGPT, Claude, Gemini, Microsoft Copilot, Midjourney.
- Hallucination — When an AI confidently generates false, misleading, or nonsensical information.
- Training data — Information fed into an AI model to teach it. Company data submitted to a public AI tool may become part of that model and could be exposed to others.
- Shadow IT — Use of unauthorised or unvetted software by staff without IT department approval.
- Jailbreaking / prompt injection — Crafting inputs to trick an AI into bypassing its safety filters or security rules.
- PII (Personally Identifiable Information) — Data that could identify an individual: names, NI numbers, addresses, phone numbers, email addresses.
- Anonymisation — The process of removing or masking sensitive details (names, prices, identifiers) from a document before sharing it with an AI tool.

- Deepfake — Synthetic media (video, audio, image) generated by AI to convincingly mimic a real person, often used for impersonation or fraud.
- API (Application Programming Interface) — A software bridge that lets two applications exchange data. Automation tools (Zapier, Power Automate, Make.com) use APIs to send data between systems.
- Human oversight — A process requiring human review or approval before an AI-influenced decision is finalised.
- BYOD (Bring Your Own Device) — The practice of staff using personal devices for work purposes.

4. Scope of policy

This policy applies to all forms of AI usage, including but not limited to:

- Web-based tools (e.g. ChatGPT, Claude, Gemini)
- Browser extensions (grammar checkers, page summarisers, AI sidebars)
- Operating system features — AI features embedded in Windows, macOS, iOS or Android (Copilot in Windows, Apple Intelligence)
- Meeting recorders — automated transcription tools
- AI features and tools pre-installed on company devices

Note: AI features bundled into apps you already use (Word, Outlook, Teams) are also covered by this policy.

5. Authorised and prohibited software

To protect business and client data, we restrict which AI tools may be used for company purposes. Using unvetted tools exposes the business to Shadow IT risks — data quietly siphoned off to train public models or stored on servers we don't control.

You may only use AI tools that have been formally approved and provisioned by IT. Expenses for unapproved AI tool subscriptions will not be reimbursed.

Refer to Appendix A for the current Authorised Software list.

IMPORTANT — Any AI tool, browser extension or plugin not listed in Appendix A is strictly prohibited.

We do not maintain a list of "banned apps" because new tools appear daily. Unless a tool has been vetted and added to the Authorised list, it cannot be used for company business.

6. Requesting new tools

We encourage innovation. To request a new AI tool be added to the Authorised list, contact Inology IT.

We will run an AI Vendor Review covering privacy, security and data governance to ensure the tool is safe for the environment in which it would operate.

7. Data classification protocol

You must classify data before entering it into any AI tool. The matrix in Part 2 of this pack is the quick reference.

If you are unsure, stop and ask. A two-minute Data Classification Review with IT is always cheaper than a breach notification.

8. Regional and industry-specific regulations

Processing data subject to specific privacy or industry regulations requires a compliance review by IT before AI use. Using AI with regulated data — without the right legal agreements in place — can lead to serious penalties.

Most standard AI tools do not automatically meet the higher standards required by GDPR, healthcare regulations, financial services rules, or legal sector compliance. Many regulations require data to remain in a specific jurisdiction, or require strict audit logs that public AI tools simply do not provide.

You are responsible for adhering to the laws and regulations relevant to your role. Contact IT or the Legal Team if you are unsure.

9. Client contract requirements

Client contracts come first. Before using AI on a client project, verify that the relevant Master Services Agreement (MSA) does not prohibit the use of AI.

If a contract forbids AI usage, you must not use it — regardless of data classification.

10. Meeting recorders and extensions

- Automated meeting recorders — Third-party AI bots (Otter.ai, Fireflies, etc.) are prohibited from joining meetings unless authorised by IT. If an unauthorised bot joins, the meeting host must remove it immediately.
- Browser extensions — Installing browser extensions that use AI features is prohibited unless deployed by IT. Many extensions request full read-access to web traffic, including private email, banking portals, and client systems.

11. Monitoring and privacy

Company-provided AI accounts are company property.

Inology IT reserves the right to audit, monitor and review all prompts, inputs and outputs generated on these accounts to ensure compliance with this policy. Company AI accounts are monitored and not private.

12. AI hallucinations and output verification

AI models hallucinate. When they do, they confidently present incorrect information as fact.

You are required to fact-check 100% of AI-generated claims against a primary source.

Never rely on AI alone for:

- Legal statutes or case law
- Mathematical calculations
- Factual citations or historical dates
- Financial figures or pricing

If you make a business decision based on AI output, you must validate the data first.

13. Intellectual property and copyright

- Copyright risks — Do not use AI to generate content that copies the distinct style of existing copyrighted works, or that deliberately reproduces protected trademarks.
- Brand representation — Be cautious when using AI to generate photos, video or audio that represents the brand. AI often introduces subtle errors (misspelled logos, anatomical anomalies) that damage professional reputation.
- Code generation — Using AI to generate software code is prohibited unless the user is qualified to audit it for security vulnerabilities and licence compatibility. AI can silently "borrow" open-source code, creating licensing issues downstream.

14. Automation workflows

Staff are prohibited from creating automated workflows (Zapier, Power Automate, Make.com) that send company data to an AI API without prior IT review.

A misconfigured automation can leak thousands of emails or files in minutes. Even when the AI tool is approved, the API connection must be secured by IT.

15. Voice cloning and deepfakes

The use of AI to clone, simulate or mimic the voice or likeness of any human being — staff, contractors, suppliers, public figures — is strictly prohibited unless explicitly authorised in writing by company leadership.

This includes text-to-speech tools trained on a specific individual's voice samples.

16. Financial verification

Deepfake technology is now used to commit wire fraud.

Any urgent request for funds (wire transfers, gift cards, invoice payments) or credential changes received via voice, video or email requires secondary verification.

If you get a call from the "CEO" asking for money, hang up. Call them back on their known internal phone number to verify the request.

17. AI-assisted decisions

AI lacks human judgement, ethical reasoning and real-world context. It hallucinates, relies on biased training data, and sounds authoritative while being completely wrong.

Relying on AI as the sole decision-maker leads to disastrous outcomes — wrongful termination claims, flawed investment strategies, regulatory fines, damaged client relationships.

A human must review and approve all AI-assisted decisions. This applies to:

- Hiring and termination decisions
- Employee performance evaluations
- Financial approvals
- Strategic business planning
- Professional advice (legal, medical, financial)

18. Transparency and AI-assisted generation

- External disclosure — If AI is used to generate any content for a deliverable (report, code module, image, article), disclose the use of AI to the client unless the contract states otherwise.
- Internal disclosure — When submitting work to a manager, flag content that was AI-assisted to ensure proper review.
- Watermarking — Where technically feasible, AI-generated images should keep their metadata or visible watermarks indicating their artificial origin.

19. Mobile and personal device usage (BYOD)

Using AI apps on personal devices to work on company tasks is prohibited.

- You may not copy company emails, files or chat logs into AI apps (e.g. the ChatGPT iOS or Android app) on a personal device.
- Mobile access to AI tools is only permitted via the company-managed Work Profile or apps installed by IT.

20. AI training requirements

Access is conditional on competence. Access to company-provisioned AI tools is granted on completion of any required AI training programmes — AI Security Awareness, AI Usability Training, Ethics modules.

The company reserves the right to revoke AI access for any member of staff who fails to complete required annual training.

21. Vendor and contractor compliance

Security standards extend to the supply chain. Contractors, vendors and freelancers are required to use company-approved AI tools when working on company data. Personal or free AI accounts by contractors are prohibited unless explicit written permission is granted by IT.

All vendors must sign this policy as part of their onboarding.

22. Termination and data retention

All prompts, inputs and outputs generated using company-provisioned AI accounts are the property of the company.

On termination of employment, IT will archive the AI account to ensure business continuity. Staff may not delete, export or transfer AI chat history to a personal account before departure.

23. Ethics and prohibited content

AI tools may not be used to generate content that violates the company's Code of Conduct or Harassment Policy.

- Prohibited uses — Generating discriminatory, sexually explicit, hateful or harassing content.
- Malicious use — Facilitating cyberattacks, creating phishing emails, or generating malicious software.

24. Security integrity (jailbreaking)

Staff are strictly prohibited from attempting to bypass the security filters, content moderation protocols or safety guardrails of any AI tool — commonly known as jailbreaking or prompt injection.

Attempting to manipulate an AI tool to ignore its safety instructions is a violation of this policy.

25. Incident reporting

If you accidentally input restricted data into an AI tool, or suspect an unauthorised bot has recorded a meeting, report it immediately. Self-reporting accidental errors is encouraged and allows for rapid containment.

To report an issue, raise an Emergency Security Ticket or call Inology IT immediately on 0161 503 3536.

The Incident Response Checklist in Part 4 of this pack walks through the first 30 minutes.

26. Policy review and updates

Due to the rapid pace of AI advancement, this is a living document.

It will be reviewed and updated by Inology IT at least every six months — or sooner if significant new AI capabilities emerge — to ensure it addresses emerging risks.

Staff will be notified of major updates to the Authorised Software list or security protocols. Continued use of company systems after an update means the new terms apply.

27. Policy enforcement and acknowledgement

These guidelines exist to protect staff, clients and the business. Mistakes happen. Wilful disregard for these safety measures — bypassing controls or sharing sensitive data — is a serious matter.

If a violation is detected, the circumstances will be investigated. Accidental errors that are self-reported are typically treated as learning opportunities. Deliberate violations or negligence will be met with appropriate disciplinary steps.

Acceptance

I have read and understood the Inology IT Artificial Intelligence Acceptable Use Policy. I agree to abide by the rules and responsibilities outlined above.

I understand that AI evolves rapidly, and it is my responsibility to seek clarification from IT or my manager whenever I am unsure about the safety of a specific tool or action. I will not assume a tool is safe just because it is publicly available.

I confirm that I understand my role in protecting company data. I agree to report any potential security risks, accidental data leaks, or suspicious activity immediately, knowing that early reporting helps protect the entire team.

Name ___

Date ___

Signature ___

PART 2

Data Classification Matrix

Before you paste anything into an AI tool, ask: what colour is this data? Green, amber, or red? This is the quick reference. Print it. Pin it.

inology I.T.

Data Classification Matrix

Use this table to classify any data before entering it into an AI tool.

Classification	AI Rating	Description	Requirements	Examples
Public data	SAFE	Information already available in the public domain.	**No restrictions.**	Marketing copy General industry research Brainstorming
Internal data	CAUTION	Internal business documents that are not confidential.	**Anonymisation required.**	Internal memos Process documentation Draft emails
Restricted data	PROHIBITED	Any data that would cause harm if leaked.	**Never input into AI.**	PII: client names, NI numbers, addresses Financials: bank accts, card numbers Credentials: passwords, API keys

If you can't tell, ask

Genuinely unsure? Stop. Don't paste. A two-minute call with Inology IT on 0161 503 3536 is always cheaper than a breach notification.

Five questions to ask yourself

- Could this data identify a person? Names, emails, phone numbers, NI numbers — anything that names a real human is PII.
- Is it covered by a client contract? Some MSAs explicitly forbid AI processing. Check before you paste.
- Would you email it to a stranger? If you wouldn't, don't paste it into a public AI tool.
- Is the AI tool on the Approved list? If not, it's prohibited — regardless of how safe the data feels.
- Does the data leave the UK or EU? Many tools default to US servers. Regulated data may not legally make that journey.

PART 3

Approved Tools Register

A living list of every AI tool the business has formally vetted and provisioned. If a tool isn't on this list, it isn't approved.

Approved Tools Register

Maintained by IT. Reviewed quarterly. Any AI tool, browser extension or plugin not on this list is prohibited for business use.

Tool	Vendor	Approved For	Data Tier	Reviewed
Microsoft Copilot (M365 Protected)	Microsoft	All staff. Productivity, summarisation, draft assistance.	Internal	May 2026
ChatGPT Team / Enterprise	OpenAI	Approved roles only. No training on inputs.	Internal	May 2026
(add new tools here after IT review)	—	—	—	—

Vetting checklist (used by IT before adding a tool)

- Data residency — Where is data processed and stored? UK/EU only, or transferred outside?
- Training opt-out — Can we guarantee company inputs are not used to train the public model?
- Authentication — SSO support? MFA enforced?
- Audit logs — Can we see who used it, when, and what they prompted?
- Data retention — How long is prompt history kept? Can we delete it on demand?
- Sub-processors — Who else does the vendor share data with?
- Compliance posture — Cyber Essentials, ISO 27001, SOC 2, GDPR DPA available?
- Termination — On contract end, is data returned and destroyed? Confirmed in writing?

How to request a new tool

- Email info@inology.co.uk with the tool name, vendor URL, and intended use case.
- IT runs the AI Vendor Review against the checklist above (typically 5 working days).
- If approved, the tool is added to this register, provisioned through SSO, and rolled out to the requested team.
- If declined, IT will explain why and (where possible) suggest a vetted alternative.

PART 4

Incident Response Checklist

What to do in the first 30 minutes after data lands somewhere it shouldn't. Speed matters. Containment now is cheaper than damage control later.

inology I.T.

Incident Response Checklist

If you suspect or have caused an AI-related data incident — restricted data pasted in error, an unauthorised meeting bot, a suspicious AI-generated request — work this checklist top to bottom. Do not delete anything until step 4.

First 5 minutes — Stop the bleeding

☐ Stop using the AI tool immediately. Don't prompt it again. Don't try to "correct" what was sent.

☐ Take a screenshot of what happened, including any conversation history and timestamps. Save to a secure location, not the AI tool itself.

☐ If a meeting bot was the issue, kick it out of the meeting and end the recording.

☐ Disconnect any automated workflow (Zapier, Power Automate, Make.com) that might still be running.

Within 15 minutes — Report it

☐ Call Inology IT on 0161 503 3536. State "AI security incident" — this triggers the priority queue.

☐ If unable to reach by phone, raise an Emergency Security Ticket and email info@inology.co.uk with subject "AI INCIDENT".

☐ Do not discuss the incident in a group chat, an email thread or a recorded meeting until IT has scoped containment.

☐ Tell your line manager you've reported it. They do not need full detail — just that an incident is in progress.

Within 30 minutes — Contain & assess

☐ IT will help identify what data was exposed, where it went, and whether it can be deleted from the tool's history.

☐ If client data is involved, check the client's MSA for incident notification timelines (often 24-72 hours under GDPR).

☐ If the tool stored credentials or API keys, rotate them immediately — don't wait for the post-mortem.

☐ Document the timeline in a single shared note: when it happened, when reported, who was contacted.

Within 24 hours — Notify & remediate

☐ IT decides whether the incident triggers GDPR notification (within 72 hours of awareness if there is a risk to individuals).

☐ Legal review for any client contracts that mandate disclosure.

☐ Update the Approved Tools register if the incident exposed a vetting gap.

☐ Add the case (anonymised) to the next quarterly AI policy review.

Self-reporting saves jobs. If you made a mistake and own up immediately, it is treated as a learning event. Mistakes that are concealed until discovered are not.

PART 5

Staff One-Pager

Print it. Pin it. Read it before you paste anything.

inology I.T.

Use AI like a sharp tool, not a lazy shortcut.

01
Classify before you paste
Public, internal or restricted? If it names a person, an account, a password or a price — assume restricted until proven otherwise.

02
Approved tools only
Microsoft Copilot or ChatGPT Team. Anything else is prohibited until IT has vetted it.

03
Fact-check everything
AI hallucinates. Verify every claim against a primary source — especially legal, financial, or factual statements.

04
No automation without IT
Zapier, Power Automate, Make.com workflows that send data to AI need IT review first. One bad rule can leak thousands of files.

05
If it asks for money, hang up
Voice cloning is real. Any urgent funds or credentials request via call, video or email needs second-channel verification.

06
Tell us when it goes wrong
Self-report incidents on 0161 503 3536. Mistakes are forgivable. Cover-ups aren't.

Questions? Inology IT · 0161 503 3536 · info@inology.co.uk

AI is the sharpest tool we've been handed in a generation. Treat it like one — respect the edge, and it'll do remarkable work for you. Be careless with it, and someone gets hurt.