ARTIFICIAL INTELLIGENCE & DATA PRIVACY POLICY

Including GDPR Compliance Framework

Version	2.0 — AI & GDPR Unified
Status	Active — For immediate distribution
Effective Date	January 2026
Next Review	January 2027
Policy Owner	Akil Benjamin
Applies To	All employees, contractors, and third parties acting on behalf of the organisation

⚠ *This is a living document and will be updated as AI and data protection law evolves.*

1. Purpose & Scope

This policy sets out a single, unified framework governing the responsible use of Artificial Intelligence (AI) and the protection of personal data across our organisation. The two subjects are deliberately combined: almost every practical AI use case involves personal data, and GDPR obligations apply directly and materially to AI adoption decisions.

Our goal is to enable our teams to work with the best available AI tools with confidence — knowing that privacy, security, and legal compliance are built in, not bolted on.

1.1 Who This Policy Applies To

- All full-time and part-time employees
- Contractors, freelancers, and consultants with access to company systems or data
- Third-party vendors and partners processing personal data on our behalf

1.2 What This Policy Covers

AI Scope	Data Privacy Scope
• Generative AI (text, image, code, audio, video) • AI-assisted analytics & decision-support tools • AI features embedded in existing software • Custom or internally developed AI models • Automated pipelines and AI agents	• All personal data collected, used, or stored • Employee, client, prospect, and supplier data • Special category data (health, biometric, etc.) • Marketing, analytics, and tracking activities • Third-party data processing arrangements

2. Key Definitions

AI Tool	Any software, platform, plugin, or API using machine learning, large language models (LLMs), or similar techniques to generate, analyse, classify, or act on content.
Personal Data	Any information relating to an identified or identifiable natural person (GDPR Article 4). Includes names, email addresses, IP addresses, location data, and online identifiers.
Special Category Data	Personal data revealing racial/ethnic origin, political opinions, religious beliefs, trade union membership, health, sex life/orientation, or biometric/genetic data (GDPR Article 9).
Data Subject	The identified or identifiable living individual to whom personal data relates.
Data Controller	The organisation that determines the purposes and means of processing personal data. This organisation acts as Controller for most processing activities.
Data Processor	A third party that processes personal data on behalf of the Controller (e.g. an AI vendor). Must be governed by a Data Processing Agreement.
Approved AI Tool	A tool that has passed our security, privacy, and legal review and is listed on the Approved AI Tools Register.
DPIA	Data Protection Impact Assessment — a structured risk assessment required before high-risk personal data processing.
DPA	Data Processing Agreement — the contractual instrument required under GDPR Article 28 with all Processors.
Human-in-the-Loop	A process requiring a human to review, verify, and take accountability for AI outputs before they are acted upon.
Lawful Basis	The legal ground under GDPR Article 6 (and Article 9 for special category data) that justifies processing personal data.
Prompt Hygiene	The practice of removing or replacing identifiable personal data before submitting inputs to AI tools.

3. Core Principles

All AI use and personal data processing must be governed by these seven principles. They are deliberately technology-neutral — they apply regardless of which tools are in use today or emerge in future.

Principle	What It Means in Practice
1. Lawfulness & Fairness	Every processing activity must have a documented lawful basis. Personal data must never be used in ways that are deceptive, discriminatory, or harmful to data subjects.
2. Purpose Limitation	Data collected for one purpose must not be repurposed without a fresh lawful basis. Define the purpose before collecting — AI exploration is not a purpose.
3. Data Minimisation	Collect and process only the minimum personal data necessary. Apply this to AI prompts: strip out PII before it is needed.
4. Accuracy	Keep personal data accurate and up to date. AI outputs based on stale or inaccurate data must not be relied upon without verification.
5. Storage Limitation	Personal data must not be retained longer than necessary. Define retention periods. Ensure AI vendor data retention aligns with our obligations.
6. Security & Integrity	All personal data must be protected by appropriate technical and organisational measures. This applies to data at rest, in transit, and within AI prompts and outputs.
7. Human Accountability	AI outputs are never final without human review. Accountability for any work product incorporating AI rests with the employee, not the tool.

> *These principles are directly drawn from the UK GDPR / EU GDPR Article 5 accountability framework and apply to all personal data processing, including AI-driven processing.*

4. Lawful Basis for Processing Personal Data

Before any personal data is processed — including via AI tools — a lawful basis under GDPR Article 6 must be identified and documented. The most common bases we rely on are:

Contractual Necessity	Processing required to perform a contract with the data subject (e.g. HR, payroll, client delivery).
Legitimate Interests	Where our interests or those of a third party are not overridden by the data subject's rights. Requires a Legitimate Interests Assessment (LIA). Commonly used for analytics, fraud prevention, and internal AI tools.
Legal Obligation	Processing required to comply with law (e.g. tax records, employment law obligations).
Consent	Freely given, specific, informed, and unambiguous indication of the data subject's wishes. Must be easy to withdraw. Required for most

	marketing and for special category data in conjunction with Article 9(2)(a).
Vital Interests	Narrowly applicable — only where processing is necessary to protect life.
Public Task	Applicable only where the organisation exercises official authority. Generally not relevant for commercial entities.

4.1 Special Category Data

Processing special category data (health, biometric, racial origin, etc.) requires BOTH a lawful basis under Article 6 AND a condition under Article 9(2). This almost always requires explicit consent or a specific legal authorisation. Special category data must never be submitted to an AI tool without written DPO sign-off.

4.2 Documenting the Basis

All processing activities involving personal data must be recorded in our Records of Processing Activities (ROPA). When AI is used to process personal data, the ROPA entry must specify the AI tool, the data categories, and the lawful basis. The DPO maintains the ROPA.

5. Data Classification & AI Use Rules

All data must be classified before being submitted to any AI tool. The table below defines the four tiers and the corresponding AI permissions:

Tier	Classification	Examples	AI Use Rule
TIER 1	Public	Published reports, press releases, marketing copy	Permitted with any approved tool.
TIER 2	Internal	Process documents, meeting notes without PII, internal comms	Approved tools only. Do not use free-tier or consumer tools.
TIER 3	Confidential	Client data (anonymised), commercial strategy, financial forecasts	Approved tools with active DPA only. Prefer private API or on-premise deployment.
TIER 4	Restricted	PII, special category data, HR records, source code, trade secrets, payment card data	PROHIBITED in external AI tools. Written DPO + Legal approval required. DPIA must be in place.

> *When in doubt, classify up. If you cannot determine the tier, treat the data as Tier 4 and consult the DPO before proceeding.*

5.1 Prompt Hygiene Rules

Even when using approved tools, employees must follow these rules before submitting any input:

- Replace real names with placeholders (e.g. “Employee A”, “Client X”).
- Remove National Insurance numbers, passport numbers, dates of birth, financial account details, and health information.
- Do not paste raw email threads containing PII without first redacting identifying details.
- Do not include passwords, API keys, or authentication credentials in any prompt.
- Review AI outputs for inadvertent PII before sharing or storing them.

6. Approved AI Tools

6.1 The Approval Process

No AI tool may be used for business purposes until it has been assessed and added to the Approved AI Tools Register maintained by IT/Legal/DPO. Assessment covers:

Assessment Area	Key Questions & Requirements
Data Processing & Location	Where is data processed and stored? Is there a DPA available? Does the vendor commit to not training on customer data?
Security Certification	Does the vendor hold ISO 27001, SOC 2 Type II, or equivalent? What encryption standards (TLS 1.2+, AES-256) are in place?
GDPR & UK GDPR Compliance	Is the vendor compliant? Are SCCs in place for non-EEA transfers? Is the vendor registered with a supervisory authority?
Access Controls	Does the tool support SSO and MFA? Are role-based permissions available? Can audit logs be exported for our records?
Model Transparency	Is the underlying model disclosed? Are known limitations and biases documented? Is there an AI usage policy from the vendor?
Data Retention & Exit	Can data be deleted on request within GDPR timescales? What is the vendor’s retention policy? Is data portability supported?
Sub-processors	Who are the vendor’s sub-processors? Are they disclosed? Are they adequately protected under our DPA?

6.2 Requesting a New Tool

Any employee may request assessment of a new AI tool by submitting a request to Akil@does.business. Requests will be reviewed within 72 hours. Using an unapproved tool for business purposes is a policy violation.

6.3 Free Tiers & Consumer Accounts

Free-tier AI tools frequently use submitted content to train their models. Employees must NEVER input company, client, or personal data into free-tier or consumer AI tools unless an explicit DPA is in place. Use only organisation-licensed accounts.

7. Data Subject Rights

GDPR grants individuals a comprehensive set of rights over their personal data. These rights apply regardless of whether the data is processed by humans or AI systems. Employees who receive a data subject rights request must forward it immediately to Akil@does.business.

Right	What It Means	Response Deadline
Right of Access (SAR)	Individuals can request a copy of all personal data held about them, including data processed by AI systems.	1 calendar month (extendable to 3 months for complex requests)
Right to Rectification	Individuals can require inaccurate personal data to be corrected, including data used as AI model inputs.	1 calendar month
Right to Erasure ('Right to be Forgotten')	Individuals can request deletion where data is no longer necessary, consent is withdrawn, or processing was unlawful.	1 calendar month
Right to Restriction	Individuals can request that processing is paused while a dispute is resolved.	Without undue delay
Right to Portability	Individuals can request their data in a structured, commonly-used, machine-readable format.	1 calendar month
Right to Object	Individuals can object to processing based on legitimate interests or for direct marketing purposes.	Immediately for marketing; otherwise 1 month
Rights re: Automated Decisions	Individuals have the right not to be subject to solely automated decisions with significant effects, unless consent or legal basis applies (GDPR Article 22).	Must be addressed before deployment

> *AI systems must not make final decisions with legal or similarly significant effects on individuals (e.g. hiring, credit, insurance) without human review. This is a hard legal requirement under GDPR Article 22.*

8. Data Retention & Deletion

Personal data must not be retained longer than is necessary for its stated purpose. This principle applies equally to data stored in AI vendor systems, logs, and training datasets.

Employee Data	Duration of employment + 6 years (employment claims limitation period). Special category HR data: review at 2 years.

Client/Prospect Data	Duration of contract + 6 years, or shorter if contractually agreed. Marketing data retained only while consent is valid.
Financial Records	7 years (HMRC requirement).
AI Prompt & Output Logs	Maximum 90 days unless required for audit or legal hold. Logs containing PII must be subject to the same retention limits as the underlying data category.
Backup Data	Maximum 12 months. Backups must be purged on schedule; they do not exempt data from erasure obligations.
CCTV / Biometric	Maximum 31 days unless required for active investigation.

Employees must not retain personal data in personal cloud storage, local drives, or AI tool chat histories beyond these limits. Use the designated records management system.

9. Security Requirements

9.1 Access Controls

- All AI tool accounts must be provisioned through IT and linked to corporate credentials where possible.
- Multi-factor authentication (MFA) is mandatory for all AI platforms.
- Shared or generic AI tool accounts are prohibited. Each user must have an individual account.
- Access must be revoked immediately upon change of role or departure.

9.2 Endpoint & Network Security

- AI tools must be accessed from company-managed devices or approved BYOD devices meeting our security baseline.
- AI tools must not be accessed from public or unsecured Wi-Fi networks without VPN.
- Browser extensions or plugins providing AI functionality must be individually approved by IT before installation.

9.3 Output Security

- AI-generated code must pass the same security review as human-authored code before deployment to any environment.
- Do not input proprietary source code, infrastructure credentials, or API keys into external AI tools.
- AI outputs incorporated into automated pipelines must be reviewed for prompt injection or adversarial content risks.

9.4 Transfers & Encryption

- All personal data in transit must be encrypted using TLS 1.2 or higher.

- Personal data at rest must be encrypted using AES-256 or equivalent.
- Transfers of personal data to AI vendors outside the UK/EEA must be governed by approved transfer mechanisms (SCCs, adequacy decision, or Binding Corporate Rules).

9.5 Incident Response

If you suspect sensitive or personal data has been shared with an AI tool without authorisation, or that a tool or system has been compromised:

- Report immediately to Akil@does.business.
- Do not attempt to recover or delete data yourself.
- Preserve records of what was submitted and when.
- Do not discuss the incident with colleagues or externally until advised.

> *Personal data breaches must be reported to the ICO (UK) within 72 hours of discovery where they are likely to result in risk to individuals' rights and freedoms. Failure to report is a regulatory offence.*

10. Acceptable & Prohibited Use

10.1 Encouraged Uses

- Drafting, editing, and proofreading internal documents (non-sensitive, no PII)
- Summarising publicly available research or long internal documents
- Writing and reviewing code (excluding credentials or proprietary algorithms)
- Generating ideas, outlines, and creative concepts
- Automating repetitive, low-risk administrative tasks with no personal data
- Learning, upskilling, and professional development

10.2 Permitted with Caution (Manager Approval Required)

- Using AI in client-facing deliverables — must be reviewed, verified, and signed off by a senior team member
- Using AI to assist in HR processes (e.g. drafting job descriptions) — HR and Legal review required
- Automated AI-generated communications sent on behalf of the organisation
- Using anonymised or pseudonymised personal data in AI tools for analysis

10.3 Prohibited Uses

- Inputting Tier 3 or Tier 4 data without explicit DPO and Legal approval
- Using AI to make unilateral employment, credit, disciplinary, or legal decisions without human review
- Processing special category data through any external AI tool without a DPIA and written DPO sign-off
- Creating deepfakes, synthetic media of real individuals, or misleading content

- Using AI to circumvent security controls, access controls, or legal obligations
- Submitting client or third-party confidential data to AI tools without contractual authorisation
- Using unapproved AI tools for any business purpose
- Misrepresenting AI-generated outputs as entirely human-produced where this would be materially misleading
- Using personal data for AI training or model fine-tuning without a documented lawful basis and DPO approval

11. Risk Tiers

Use this framework to calibrate the level of oversight and controls required. Both AI risk and data protection risk are reflected in these tiers:

Tier	Description & Required Controls
LOW	Internal productivity tasks with no personal data. Standard policy compliance sufficient. Example: drafting meeting agendas, summarising public reports.
MEDIUM	Client-facing outputs, automated communications, or tasks involving Tier 2/3 data. Manager approval + human review before use. Data minimisation applies.
HIGH	Processing personal data (Tier 3/4), HR or legal decisions, AI integrated into core business processes. DPO review, DPIA, legal sign-off, and ongoing audit required.
CRITICAL	Fully autonomous AI agents, real-time decisions with significant effects on individuals, or regulated activities. Board/executive sign-off, formal AI governance programme, and supervisory authority notification may be required.

12. Governance & Roles

Policy Owner — Akil Benjamin	Maintains and updates this policy. Chair policy review. Final escalation point for AI and data privacy queries. Contact: Akil@does.business
Data Protection Officer (DPO)	Advises on GDPR obligations. Approves high/critical risk AI deployments. Conducts and reviews DPIAs. Handles data subject rights requests. Maintains the ROPA. Manages ICO notifications.
IT / Security	Maintains the Approved AI Tools Register. Reviews tool security posture. Manages access provisioning, MFA enforcement, and audit log collection.
Legal	Reviews vendor agreements, DPAs, and SCCs. Signs off on automated decision-making use cases. Advises on lawful basis and regulatory changes.
Line Managers	Ensure team compliance with this policy. Approve medium-risk AI use cases. Escalate queries to DPO or IT appropriately.
All Employees	Read, understand, and comply with this policy. Complete mandatory training. Report breaches or suspected incidents immediately. Apply prompt hygiene at all times.

12.1 Policy Review Cadence

- Formal annual review by Policy Owner, DPO, IT, and Legal (next review: January 2027).
- Ad hoc reviews triggered by: significant AI regulatory changes (EU AI Act, ICO guidance), major vendor changes, data breaches, or significant new AI capabilities.
- All employees notified of material changes within 3 days of update.

12.2 Records of Processing Activities (ROPA)

The DPO maintains a ROPA documenting all personal data processing activities as required by GDPR Article 30. Any new processing activity — including deployment of an AI tool that processes personal data — must be registered in the ROPA before going live.

13. Training & Awareness

- All employees must complete the combined AI Literacy & Data Privacy training within 30 days of this policy's effective date and annually thereafter.
- New joiners must complete training within 14 days of their start date as part of onboarding.
- Employees operating AI tools for medium or high-risk tasks must complete role-specific advanced training before being granted elevated access.
- Training completion is tracked and is a condition of continued AI tool access.
- The DPO will communicate regulatory updates (e.g. new ICO guidance, EU AI Act milestones) to all staff as they arise.

14. Compliance & Consequences

Compliance with this policy is mandatory. Breaches will be handled in line with our Disciplinary Policy and may result in:

- Immediate suspension of AI tool access and/or system access
- Formal disciplinary action up to and including dismissal for gross misconduct
- Personal civil or criminal liability where data protection law is wilfully breached
- Mandatory reporting to the ICO or relevant supervisory authority
- Fines of up to £17.5 million or 4% of global annual turnover (whichever is higher) in the most serious GDPR breach cases

> *We operate a speak-up culture. If you are unsure whether a use case is permitted, ask Akil@does.business before proceeding. Genuine, good-faith queries will never be penalised.*

15. Amendment Log

Version	Date	Author	Summary of Changes
1.0	Jan 2026	Akil Benjamin	Initial AI Usage Policy created.

Version	Date	Author	Summary of Changes
2.0	Jan 2026	Akil Benjamin	AI & GDPR policies unified into a single document. Full GDPR framework added: lawful basis, data subject rights, retention schedules, ROPA obligations, breach notification.

This is a living document. AI technology and data protection law are both evolving — and so will this policy.

Questions? Contact Akil@does.business or speak to your line manager.