Artificial Intelligence (AI) Policy

Responsible Director:	**Director of Finance and Business**
Review Date:	**July 2026**
Next Review Date:	**July 2028**
Version:	**1.0**

Contents

1. Policy Statement

People First (PF) and its partner organisations recognise that Artificial Intelligence (AI) can improve organisational effectiveness, reduce administrative burden, support innovation, enhance reporting and business intelligence, and improve the quality of routine documentation and meeting outputs. The organisation is committed to using AI responsibly, lawfully, ethically and sustainably in a way that supports its charitable purposes and organisational values.

For the purpose of this policy, references to People First (PF) apply equally to partner organisations operating under this governance framework where AI systems are used on behalf of the organisation.

This policy applies to trustees, directors, employees, contractors, volunteers, sessional staff and any other person using AI on behalf of the organisation.

The organisation's current strategic objectives for AI are to:

- Improve organisational reporting, insight and business intelligence.
- Reduce administrative effort and duplication.
- Improve the quality, consistency and efficiency of routine documentation and meeting outputs.
- Support accessibility, inclusion and digital participation where appropriate.
- Enable employees to spend more time supporting beneficiaries and delivering organisational objectives.
- Achieve these benefits without compromising information security, confidentiality, safeguarding responsibilities, legal compliance or public trust.

The organisation recognises that AI is primarily an information governance challenge rather than a technology challenge. Users remain responsible for all information entered into, accessed through or generated by AI systems. Approval of an AI platform does not remove obligations relating to confidentiality, data protection, safeguarding, professional judgement or organisational policy.

The organisation's primary AI-related risk is the inappropriate disclosure of information. All AI use must therefore prioritise the protection of personal, confidential, safeguarding, commercial and organisational information.

This policy should be read alongside all information governance, IT, data protection, cyber security and safeguarding policies.

AI Golden Rules

1. Use approved AI tools only.
2. Do not enter personal, confidential, safeguarding or sensitive information into public or unapproved AI tools.
3. Check all AI outputs before using or sharing them.
4. Use AI to support judgement, not replace it.

5. Report concerns, incidents or unsafe outputs immediately.
6. Ask for advice before using a new AI tool or AI-enabled feature.

2. Principles

2.1 Information Governance First

Information governance takes precedence over convenience, efficiency or automation.

All users must:

- Understand the difference between approved organisational AI systems and public AI services.
- Apply organisational information classification requirements before using AI.
- Follow data minimisation principles.
- Use only approved AI systems for organisational purposes.
- Ensure AI-generated outputs are reviewed by a human before use.
- Remain accountable for the accuracy and appropriateness of all AI-assisted work.

2.2 Human Oversight and Accountability

AI may support decision-making but must not replace human judgement.

Humans remain accountable for:

- Decisions affecting beneficiaries.
- Safeguarding decisions.
- Employment decisions.
- Financial decisions.
- Strategic decisions.
- Legal or regulatory decisions.

No significant decision affecting an individual's rights, welfare, employment, eligibility or safeguarding status shall be made solely by an AI system.

2.3 Security and Confidentiality

The organisation will use AI systems that support security, resilience and appropriate protection of organisational information.

All users must protect confidential, personal and sensitive information and ensure it is processed only within approved and authorised environments.

2.4 Responsible Innovation

The organisation encourages responsible experimentation and innovation using approved AI technologies within defined governance boundaries.

Innovation should:

- Deliver organisational value.
- Improve productivity.
- Reduce administrative burden.
- Improve accessibility and inclusion.
- Support organisational sustainability.
- Remain compliant with legal and ethical obligations.

2.5 Environmental Stewardship

The organisation recognises that AI technologies consume energy, water and computing resources.

When selecting and deploying AI technologies, the organisation will seek to balance innovation and productivity benefits with environmental responsibility and will consider environmental impacts alongside cost, performance, security and compliance requirements.

3. Law and Legislation

3.1 Regulatory Approach

The organisation will comply with all applicable legislation relating to AI, privacy, data protection, cyber security and information governance.

Where legal requirements are unclear or evolving, the organisation will adopt recognised good practice and proportionate risk management.

3.2 EU AI Act

The organisation adopts the principles of the EU AI Act as a recognised governance framework and best practice standard.

AI systems will be assessed using a risk-based approach and categorised according to organisational risk, legal obligations and potential impact on individuals.

Where AI systems are considered high risk, additional governance, approval and monitoring requirements will apply before deployment.

3.3 UK GDPR and Data Protection Act 2018

The organisation applies GDPR principles to all AI-related processing involving personal data.

This includes:

- Lawfulness, fairness and transparency.
- Purpose limitation.
- Data minimisation.
- Accuracy.

- Storage limitation.
- Integrity and confidentiality.
- Accountability.

Any AI-related data breach shall be managed through existing data breach procedures.

3.4 UK Government Artificial Intelligence Playbook

The organisation will adopt relevant principles from the UK Government Artificial Intelligence Playbook, including:

- Responsible use.
- Human oversight.
- Security.
- Transparency.
- Accountability.
- Lifecycle management.
- Continuous improvement.

3.5 Charity Governance

Trustees and Directors must ensure AI use furthers the organisation's charitable purposes, protects beneficiaries, manages risks appropriately and supports the principles of good charity governance.

AI must only be used where it supports the organisation’s charitable purposes, public benefit, service quality, safeguarding responsibilities and responsible use of charitable resources. AI must not be used in a way that undermines beneficiary trust, creates unfair barriers to services or weakens human relationships with people who use the organisation’s services.

4. Information Classification and AI

Before using any AI system, users must consider the classification of the information involved. The classification determines whether AI may be used, which system may be used, and what safeguards are required.

AI must not be used simply because it is convenient. Users must apply the same judgement they would apply when sharing information by email, storing information in SharePoint, or discussing information with a third party.

4.1 General principles

All users must:

- Check whether the information is public, internal, confidential, restricted, special category, safeguarding-related, commercially sensitive or otherwise high risk.

- Use the minimum amount of information necessary to complete the task.
- Avoid entering personal data unless there is a clear organisational need, an approved system is being used, and the use is consistent with this policy.
- Anonymise or summarise information wherever possible.
- Never use AI to bypass information governance, confidentiality, safeguarding, cyber security or professional judgement requirements.
- Review all AI outputs before relying on them, sharing them or placing them on an organisational record.
- Ask their manager, the Data Protection Officer, the Information Security Lead or the Director of Finance and Business if they are unsure whether information can be used with AI.

4.2 Classification rules

Information classification	Examples	AI use position
Public information	Published website content, public reports, general research, publicly available guidance	May be used with approved AI systems. Public or consumer AI may only be used where no organisational, personal, confidential or sensitive information is included.
Internal organisational information	Internal processes, draft policies, meeting agendas, general planning documents, non-sensitive management information	May be used in approved organisational AI systems such as Microsoft 365 Copilot, subject to user permissions, human review and this policy.
Confidential organisational information	Non-public strategy, Board papers, internal financial information, contracts, procurement information, business planning, commercially sensitive information	May only be used in approved organisational AI systems where there is a clear business need and appropriate access permissions. Users must apply data minimisation and must not use public or consumer AI tools.
Personal data	Information relating to employees, volunteers, contractors, beneficiaries, donors, partners or members of the public	Must only be used in approved organisational AI systems where there is a clear lawful purpose, the minimum necessary information is used,

		and the use is consistent with privacy notices and data protection requirements. Personal data must not be entered into public or unapproved AI systems.
Special category or highly sensitive data	Health information, disability information, ethnicity, religion, trade union membership, sexual orientation, biometric data, criminal offence information, disciplinary records, grievance records, investigations, safeguarding information	Must not be entered into AI systems unless specifically authorised through an approved governance process supported by risk assessment, Data Protection Officer review and appropriate safeguards. Public or consumer AI tools must never be used for this information.
Safeguarding information	Safeguarding concerns, referrals, risk assessments, case notes, disclosures, allegations, or information relating to children or vulnerable adults	Must not be entered into AI systems unless specifically authorised through an approved governance process. AI must never be used to make safeguarding decisions or replace professional judgement.
Authentication, security or system information	Passwords, access tokens, security configurations, vulnerability information, network information, system credentials	Must not be entered into AI systems unless expressly authorised by the Information Security Lead for a defined security purpose using an approved system.
Third-party confidential information	Supplier confidential information, partner data, contractually restricted material, information received under confidentiality arrangements	Must only be used in approved organisational AI systems where the organisation has the right to use the information in this way and where contractual requirements permit it.
AI-generated outputs	Draft text, summaries, analysis, meeting notes, reports, recommendations, code, images or other outputs created with AI support	Must be checked for accuracy, bias, completeness, appropriateness and compliance before use. Users remain accountable for the final output.

4.3 Public and consumer AI services

Public or consumer AI services must not be used for organisational information unless explicitly approved through the governance process.

Staff must not enter any of the following into public or consumer AI services:

- Personal data.
- Confidential organisational information.
- Commercially sensitive information.
- Financial information that is not already public.
- HR, disciplinary, grievance, recruitment or investigation information.
- Safeguarding information.
- Beneficiary, service user or donor information.
- Security, password, access or system information.
- Information received from third parties under confidentiality obligations.

4.4 Microsoft 365 Copilot and organisational information

Microsoft 365 Copilot may be used for approved organisational purposes where the information is already available to the user through their normal Microsoft 365 permissions and where the use is consistent with this policy.

Users must remember that Microsoft 365 Copilot can surface information based on existing access permissions. If information is incorrectly shared, stored in the wrong location or accessible to the wrong users, Copilot may make that information easier to find. The correct control is to manage permissions, information storage and sharing practices properly.

4.5 Labelling and transparency

Where AI has materially contributed to a document, report, policy, communication or analysis, users should be transparent about its use where appropriate. The level of disclosure should be proportionate to the context, audience and risk.

AI-generated or AI-assisted content must not be presented as final organisational advice, professional judgement or approved policy unless it has been reviewed and approved by the appropriate person.

5. Governance and Accountability

Board of Trustees

The Board of Trustees is responsible for approving the organisation's AI risk appetite and ensuring that AI use remains proportionate to the organisation's charitable objectives, regulatory duties and responsibilities to beneficiaries.

Executive Directors

Executive Directors share responsibility for ensuring AI is used appropriately within their respective functions and that risks, incidents or concerns are escalated through the organisation's governance arrangements.

Director of Finance and Business

The Director of Finance and Business owns this policy on behalf of the organisation, coordinates AI governance activity and ensures the AI inventory, assurance reporting and policy review process remain current.

Data Protection Officer

The Data Protection Officer advises on DPIAs, lawful processing, data subject rights, privacy compliance and the safeguards needed where AI use involves personal data.

Information Security Lead

The Information Security Lead is responsible for security controls, access management, incident response, supplier assurance and technical monitoring of approved AI systems.

Service Owners

Service owners are responsible for ensuring AI use within their service, project or business area is justified, proportionate, tested and subject to appropriate human oversight. Their responsibilities include:

- Owning the business case for AI use within their area.
- Confirming the purpose, expected benefit and proportionality of AI use.
- Ensuring AI use aligns with this policy, information governance requirements and operational need.
- Ensuring appropriate testing is undertaken before AI is relied upon operationally.
- Defining success measures and monitoring whether expected benefits are being achieved.
- Ensuring human review and professional judgement remain in place.
- Identifying and reporting risks, issues, incidents and near misses.
- Supporting completion of DPIAs, supplier checks and AI inventory records where required.
- Ensuring staff within their area understand permitted and prohibited AI use.

All Users

All users are responsible for following this policy, protecting organisational information, using approved AI systems only, reporting incidents or concerns and completing required AI training.

6. AI Inventory and Classification

The organisation will maintain an AI inventory so that approved, restricted, prohibited and retired AI systems are visible and subject to appropriate governance.

The inventory will record approved AI systems, their business purpose, system owner, supplier, risk classification, data categories processed, review date, environmental considerations where known and lifecycle status.

All AI systems must be recorded before organisational deployment and reviewed when their purpose, supplier, risk profile or lifecycle status changes.

7. Lifecycle Controls

Planning and Design

Before adopting a new AI system, the organisation will define the business purpose, assess expected benefits, consider non-AI alternatives, evaluate risks, consider ethical implications and assess information governance requirements.

Planning evidence should include a proportionate business case, risk assessment and, where personal data is involved, consideration of whether a DPIA is required.

Data Controls

The organisation will use data lawfully and ethically, apply data minimisation, protect personal information and ensure information remains within approved environments.

Where an AI use case involves personal data, the Data Protection Officer shall advise on whether a DPIA, privacy notice update, lawful basis review or additional safeguard is required.

Deployment Controls

The organisation will restrict use to approved AI systems, provide user training, maintain governance records and monitor usage and emerging risks.

Deployment evidence should include user guidance, training requirements, approval records and confirmation that the AI system has been added to the AI inventory.

Monitoring Controls

The organisation will review effectiveness, productivity benefits, incidents, near misses, compliance and emerging risks to ensure AI remains safe, useful and proportionate.

Monitoring evidence should include usage information where available, incidents, near misses, user feedback, control weaknesses and benefits realised.

Decommissioning Controls

When retiring AI systems, the organisation will remove access, manage data appropriately, review lessons learned and update the AI inventory.

Decommissioning evidence should include access removal, data retention or disposal decisions, supplier exit arrangements and lessons learned.

8. Approved AI Tools

Only AI systems approved by the organisation may be used for organisational purposes.

An approved organisational AI system is an AI system that has been reviewed, authorised, recorded in the AI inventory and is subject to organisational controls, including access management, user guidance, monitoring where available and periodic review.

This policy applies to standalone AI tools and AI-enabled features embedded within existing systems.

Approved systems currently include:

- Microsoft 365 Copilot products licensed to and administered within the organisation's Microsoft 365 tenant.
- Microsoft Copilot consumer/web only for general, non-confidential tasks where no organisational, personal, confidential, restricted, safeguarding, HR, financial, commercially sensitive or otherwise sensitive information is entered.
- Other AI systems that have been specifically approved through organisational governance processes and added to the AI inventory.

Microsoft 365 Copilot may be used for approved organisational activities subject to this policy and existing information governance requirements. Prompts, responses and data accessed through Microsoft 365 Copilot must remain within the organisation's Microsoft 365 tenant and must not be used to train public AI models.

Public or consumer AI tools, including ChatGPT, Gemini and similar services, are prohibited for organisational use unless separately approved through the governance process following a business case, information governance review and risk assessment.

Users must not assume that a tool is approved simply because it has Microsoft branding. Microsoft 365 Copilot and consumer/web Copilot are subject to different controls and different permitted uses.

Plaud AI is not an approved organisational AI tool under this policy. It must not be used for organisational recording, transcription, note-taking or meeting summary purposes unless it is subject to a future approval decision and added to the AI inventory.

Users must not input special category data, safeguarding information, confidential beneficiary details, HR information, financial records or other restricted information into any AI system unless this is expressly authorised under this policy and through the appropriate governance process.

Approval of a platform does not remove the responsibility of users to apply professional judgement, data minimisation, confidentiality and information governance controls.

Microsoft Purview DSPM for AI may provide visibility of sensitivity labels, where configured, and can support monitoring, risk review and governance. It may help identify potential AI-related risks such as sensitive information exposure, oversharing risks, risky interactions and usage trends.

However, the organisation's current licence plan does not provide sensitivity label enforcement, full DLP configuration, detailed external AI usage monitoring or detailed data transfer monitoring to the level required to rely on these as complete technical controls.

The organisation's current controls therefore remain:

- User training.
- Clear policy requirements.
- Approved tool lists.
- Human review and professional judgement.
- Access and permission management.
- Available Microsoft Purview/DSPM monitoring.
- Management oversight and assurance reporting.

The organisation will keep its technical control environment under review and will consider whether additional Microsoft Purview, DLP, sensitivity labelling or monitoring capabilities are required as AI use develops.

Before wider deployment of Microsoft 365 Copilot, the organisation should review SharePoint, Teams and OneDrive permissions to reduce oversharing risk and ensure users only have access to information required for their role.

9. Supplier and Third-Party Management

Before adopting AI systems, the organisation will assess whether suppliers can meet legal, security, privacy, operational and environmental expectations.

Supplier due diligence will consider security arrangements, data residency, regulatory compliance, sub-processors, previous incidents, environmental impacts where available and contractual protections.

Contracts should include appropriate provisions for data processing, information security, audit rights, service levels and data transfer requirements.

New AI tools, AI-enabled features or material new AI use cases must not be adopted until they have been reviewed by the relevant service owner, Data Protection Officer, Information Security Lead and Director of Finance and Business. Approval shall be based on business need, risk, data protection, cyber security, supplier assurance, cost, environmental impact and whether the tool is added to the AI inventory.

Approval of AI tools must also comply with the organisation's procurement procedures, budget controls and delegated authority levels. No AI tool may be purchased, trialled or renewed outside approved budget and authorisation arrangements.

Supplier due diligence should include, where proportionate, review of the supplier's legal role, security arrangements, data residency, sub-processors, previous incidents, environmental impact, data processing terms, international transfer arrangements, service levels and audit rights.

10. Safeguarding and High-Risk Information

Safeguarding information must not be entered into AI systems unless specifically authorised through an approved governance process.

Any such authorisation must be supported by risk assessment, legal review, information governance review and appropriate safeguards.

AI must never be used to make safeguarding decisions or replace professional judgement.

Employment relations, disciplinary matters, grievances, investigations and recruitment assessments must not be entered into unauthorised AI systems.

11. Data Subjects and Beneficiaries – Rights and Safeguards

The organisation will respect data subject and beneficiary rights, including rights to information, access, rectification, restriction, objection, portability and human review where appropriate.

The organisation will take additional care when processing information relating to children and vulnerable adults.

12. Environmental and Ethical Considerations

The organisation will consider environmental and ethical factors alongside cost, performance, security, accessibility and compliance when selecting and using AI systems.

This includes considering environmental impacts, preferring efficient and lower-impact technologies where appropriate, encouraging supplier disclosure, using AI responsibly and ethically, monitoring potential bias and promoting accessibility and inclusion.

13. AI Assurance and Reporting

Training and AI Literacy

Training will cover approved and prohibited AI tools, safe prompting, information classification, data minimisation, personal and special category data, safeguarding, bias, accessibility, human judgement, output checking, environmental and ethical considerations, and incident escalation.

AI training will be included in induction for relevant users and refreshed at least annually. Additional targeted training may be required for users with Microsoft 365 Copilot licences, service owners, managers, system owners, the Data Protection Officer, the Information Security Lead and staff involved in procurement, reporting, safeguarding, HR, finance or service delivery.

Review and Assurance

The Director of Finance and Business will coordinate annual review of this policy and related AI governance arrangements.

The AI inventory, approved tools list, risk register entries, training completion, incidents, near misses, DPIAs, supplier assurance records and available Microsoft 365 Copilot usage reports will be reviewed at least annually and more frequently where risk, legislation, organisational change or technology changes require it.

Governance records relating to AI approvals, DPIAs, supplier checks, incidents, training and assurance reporting shall be retained in accordance with organisational retention requirements. Where system reporting is limited by licence restrictions or retention periods, this limitation shall be recorded and considered as part of the annual assurance review.

The Board shall receive an annual AI assurance report covering:

- Approved AI systems and any changes to the AI inventory.
- New AI use cases approved or rejected.
- Significant risks, incidents and near misses.
- Any data protection, safeguarding, cyber security or confidentiality issues relating to AI.
- Training completion and AI literacy activity.
- Policy compliance and exceptions.

- Benefits realised, including productivity, reporting, accessibility and administrative improvements.
- Environmental and ethical considerations where known.
- Planned future developments.
- Any recommended changes to policy, tools, licences, controls or governance arrangements.

Significant AI-related incidents shall be escalated outside the normal reporting cycle where appropriate.

14. Risk Management

AI-related risks shall be recorded and managed through the organisational risk management framework, with incidents investigated, reviewed and escalated where required.

AI incidents include information disclosure or data leakage, unauthorised AI use, security breaches, biased or discriminatory outputs, significant inaccuracies, regulatory non-compliance and reputational damage.

Suspected AI-related incidents, unsafe outputs, unauthorised tool use, inappropriate disclosure, data breaches, safeguarding concerns or security concerns must be reported immediately to the user's manager and escalated to the Data Protection Officer, Information Security Lead or Director of Finance and Business as appropriate.

15. Success Measures

The organisation will assess AI implementation against objectives including no significant AI-related data breaches, improved reporting and insight, reduced administrative burden, improved productivity and efficiency, appropriate legal and regulatory compliance, and responsible and sustainable use of AI. These objectives shall be reviewed annually as part of the formal policy review process.

Appendix 1 – AI Inventory and Classification Register

Platform	Status	Permitted use	Key controls and restrictions
Microsoft 365 Copilot	Approved	Workplace productivity using organisational Microsoft 365 data, including drafting, summarising, meeting preparation, document support, reporting assistance and administrative support.	Use is subject to Microsoft 365 permissions, user training, human review, information governance controls and ongoing monitoring where available. System owner: Director of Finance and Business / Microsoft 365 administrator. Lifecycle stage: deployment and monitoring.
Microsoft Copilot consumer/web	Restricted use only	General non-confidential assistance, public research, idea generation and writing support.	Must not be used with organisational, personal, confidential, restricted, safeguarding, HR, financial, commercially sensitive or otherwise sensitive information. Individual users remain responsible for compliance.
ChatGPT	Prohibited unless separately approved	Not approved for organisational information or work purposes unless a specific approval is granted through governance.	Must not be used for organisational information unless formally assessed, approved, added to the AI inventory and subject to appropriate controls.
Gemini	Prohibited unless separately approved	Not approved for organisational information or work purposes unless a specific approval is granted through governance.	Must not be used for organisational information unless formally assessed, approved, added to the AI inventory and subject to appropriate controls.
Any other AI platform not listed above	Prohibited unless approved	Any AI tool, app, website, browser extension, embedded assistant or AI-enabled	Must not be used until assessed, approved, added to the AI inventory and subject to appropriate controls.

		service not included in this register.	

Any AI platform not shown as approved in this appendix is prohibited for organisational use unless formally approved through the governance process. The risk levels in this appendix reflect the organisation's internal risk management position and are not a substitute for legal classification under the EU AI Act or any other legislation.

Appendix 2 – Acceptable AI Use for Staff

Allowed use	Not allowed
Using Microsoft 365 Copilot to help draft emails, documents or reports, with human review before use.	Using unapproved AI tools, apps, websites or browser extensions for organisational purposes.
Using approved AI tools to improve spelling, grammar, tone or structure of work documents.	Entering personal data about service users, colleagues, volunteers, donors, partners or members of the public into public or unapproved AI tools.
Using approved AI tools to summarise non-sensitive information, meeting agendas, action notes or general planning information.	Entering special category or highly sensitive information, including health, safeguarding, financial, disciplinary, grievance, investigation or HR information, unless specifically authorised through governance.
Using AI to generate ideas, outlines, checklists, prompts or first drafts to support work.	Using AI to make or support decisions about people, including safeguarding, HR, recruitment, disciplinary, grievance, eligibility, service access or performance decisions.
Using Microsoft 365 Copilot for organisational information where the information is already accessible to the user through normal Microsoft 365 permissions and its use is consistent with this policy.	Uploading or pasting confidential, restricted, commercially sensitive or organisationally sensitive information into public or unapproved AI tools.
Anonymising or summarising information before using AI, where appropriate and approved.	Sharing AI accounts, passwords, licences or credentials, or using someone else's account.
Checking, editing and taking responsibility for all AI-generated outputs before using or sharing them.	Copying AI outputs directly into final work without checking for accuracy, bias, completeness, tone and appropriateness.

Asking your manager, the Data Protection Officer, the Information Security Lead or the Director of Finance and Business if you are unsure whether AI use is appropriate.	Attempting to bypass security controls, restrictions, monitoring or governance requirements relating to AI use.
Using approved AI tools to support accessibility, inclusion, plain English, translation drafts or alternative formats where the information used is appropriate.	Using AI-generated content in external communications, reports, advice, policies or records without appropriate human review and approval.
Using approved AI tools to support reporting, business intelligence, productivity and administrative efficiency.	

All users remain accountable for the accuracy, legality, confidentiality and appropriateness of AI-assisted work.