BrisDoc Patient care by people who care

Proudly employee owned.

AI Governance Policy

Version:	Owner:	Created:
1.0	Jon Rex (Head of DDaT)	23/10/2026
Published:	**Approving Director:**	**Next Review**
23/11/2026	Lewis Scandle	31/10/2027

Contents

Introduction....3
1. Scope and Purpose....6
Purpose and Expected Benefits....7
2. Legal Basis....7
3. Data Minimisation and Anonymisation....7
4. Fairness and Transparency....7
5. Healthcare Applications of AI....8
Induction Training....9
Annual Refresher Training....10
Role-Specific Training....10
Prohibition and Responsible Use Training....10
Ongoing Communication....10
Monitoring and Audit Criteria....11
Version Control....12

Introduction

Artificial intelligence (AI) has the potential to improve patient care, streamline administrative processes, and strengthen operational efficiency across BrisDoc. These benefits must be realised through the ethical, safe, lawful, and responsible use of AI. This policy establishes the framework for approving, implementing, managing, and overseeing AI systems throughout the organisation.

Approved and Prohibited Uses of AI

What Is Sanctioned

- Approved pilots, such as the use of Heidi within Practice Services.

A limited number of BrisDoc staff may be granted access to specified AI tools, such as ChatGPT, where there is a clearly defined business need and an approved, low-risk use case. Access will be controlled to support safe, lawful, transparent, and accountable use.

Authorised user arrangements

- **Identification:** Service/Function Leads nominate staff who require access for specific, justified purposes.
- **Approval:** The Information Governance Lead approves access, with oversight from the SIRO and Executive Sponsor.
- **Training:** Users must complete required AI governance and responsible-use training before access is granted.
- **Access control:** Named, individual accounts are issued; shared or generic accounts are not permitted.
- **Logging:** Digital and Cyber teams maintain audit logs of authorised user activity, with routine and incident-based review.
- **Review:** The authorised user list is reviewed at least every six months to ensure ongoing need and compliance.
- **Removal:** Access is removed immediately if the user changes role, no longer requires access, or breaches policy.
- **Responsibilities:** Users must stay within their approved use-case and must not input patient, staff, or sensitive data into public AI tools. Misuse may result in disciplinary action.

Unsanctioned uses

- Broad or routine use of conversational AI tools (e.g., ChatGPT, Copilot, or similar) in clinical settings.
- Use of generative AI across BrisDoc without formal approval, governance oversight, or defined use-cases.

- Any informal use that involves patient information, operationally sensitive data, or tasks that could impact clinical decision-making or service safety.
- Uploading patient, staff or operationally sensitive information into public AI tools is prohibited unless an approved, BrisDoc-controlled route exists.

Key risks

Data Protection & Information Governance Risks

Potential breaches if patient or organisational data is entered into external systems without safeguards.

Clinical Safety Risks

Tools may produce inaccurate or incomplete outputs, influencing decision-making.

Operational Risks

Over-reliance on technology that could be unavailable, unreliable, or difficult to audit.

Reputational Risks

Misuse or visible errors could harm BrisDoc's reputation, especially in a healthcare context where trust is critical.

Ethical & Equity Concerns

Generative AI tools may introduce or amplify bias.

Compliance Risks

Unsanctioned use may lead to non-compliance with NHS standards, regulatory expectations, or contractual obligations.

Duties, Accountabilities and Responsibilities

The following roles collectively provide BrisDoc's governance and assurance framework for AI.

Data Protection Officer (DPO)

- Support the IG Lead to ensure compliance with data protection regulations and best practices relating to AI at BrisDoc.
- Provide guidance on data privacy concerns associated with AI systems.
- Provide assurance and advice on Data Protection Impact Assessments (DPIAs) for AI-related projects.
- Act as the point of contact for data subjects and supervisory authorities regarding data protection issues related to AI technologies.

Senior Information Risk Owner (SIRO)

- Take responsibility for overseeing the governance and management of information risks associated with AI systems.
- Ensure appropriate risk management processes, controls, and policies are in place.
- Collaborate with key stakeholders to address potential risks and mitigate negative impacts of AI implementation.

- Provide strategic direction to promote the responsible use of AI technologies across BrisDoc.

Caldicott Guardian (CG)

- Ensure the processing of data complies with the Caldicott Principles, focusing on patient confidentiality.
- Ensure confidential patient information is processed legally, ethically, and appropriately.
- Provide advice and guidance to staff on the ethical use of AI in clinical settings.

Information Governance Lead

- Review and approve DPIAs for AI projects.
- Serve as the main point of contact for staff queries or concerns relating to AI technologies.
- Provide guidance on data protection in the context of AI.
- Ensure AI implementation aligns with current data protection legislation and standards.

Clinical Safety Officer

- Assess the clinical safety risks associated with AI systems used within patient care settings.
- Collaborate with relevant stakeholders to establish safety protocols and guidelines for AI deployment.
- Monitor and evaluate the performance and safety of AI systems.
- Investigate and address any incidents or concerns related to the clinical safety of AI technologies.

Executive Sponsor

- Accountable for ensuring AI use aligns with organisational strategy, approves business cases, and ensures service readiness.

Digital, Technical and Business Intelligence Teams

- Assist with the implementation, integration, and maintenance of AI systems across BrisDoc.
- Ensure AI systems are properly configured, secure, and compatible with existing IT infrastructure and conduct regular security patching and vulnerability management in line with NHS Cyber Security standards (NHS Digital Cyber Assessment Framework).
- Collaborate with vendors and other stakeholders to address technical issues and provide ongoing technical support for AI systems.

Learning and Development (L&D) Function

- Ensure policy compliance through contracts, job descriptions, and performance management.
- Lead consultation and change management when AI impacts services.
- Support disciplinary processes for policy breaches.
- Plan recruitment and workforce skills for AI literacy, ensuring diversity and inclusion.

- Deliver mandatory AI governance training tailored to staff roles.
- Maintain compliance records and provide refresher courses.
- Promote awareness and responsible AI use through campaigns and briefings.
- Create feedback loops to improve training and address emerging risks.

Employees and Authorised Users

- Use AI systems in accordance with established protocols and guidelines.
- Provide feedback on the effectiveness, usability, and impact of AI technologies in practice.
- Report any incidents or concerns related to AI system performance or safety.
- Familiarise themselves with and adhere to BrisDoc's Information Governance and Security policies, protocols, and guidelines.
- Report any concerns or issues regarding AI systems to the Information Governance Lead.

Policy Statement

1. Scope and Purpose

This policy applies to BrisDoc and to all employees, contractors, temporary and honorary workers, secondees, volunteers, and other individuals working for or on behalf of the organisation who can access BrisDoc systems or data. It covers every department and service using AI, regardless of scale or purpose, including systems developed internally and products or services procured from external providers.

The policy establishes clear requirements for the development, procurement, deployment, use, and monitoring of AI systems. Its purpose is to protect personal data, maintain ethical and clinical standards, and manage information, safety, operational, and reputational risks while enabling appropriate innovation. AI technologies, including machine learning and natural language processing, may support research, improve healthcare outcomes, and enable more efficient use of resources. Their use must comply with applicable law and regulation, respect patient rights, and sustain the trust of patients, staff, and stakeholders.

This policy sets out the principles, responsibilities, and controls that govern the responsible use of AI at BrisDoc, including requirements for privacy, transparency, accountability, human oversight, and ongoing monitoring. It is intended to help BrisDoc realise the benefits of AI while identifying and reducing foreseeable harm. The policy will be reviewed and updated as technology, legislation, national guidance, and recognised good practice evolve.

BrisDoc's use of AI will focus on:

- Consultation and back-office functions (e.g., scheduling, triage support, documentation).
- Not product development or academic research.
- Ensuring compliance with NHS and ICO guidance for safe, ethical, and lawful AI use.

Purpose and Expected Benefits

Before any AI technology is implemented, its intended purpose, anticipated benefits, affected users, and foreseeable risks must be clearly defined and documented. This includes:

- Understanding why AI is being used.
- Identifying the value, it will bring to patients, staff, and organisational efficiency.
- Ensuring the purpose aligns with BrisDoc's mission of delivering safe, high-quality, and patient-centred care.

2. Legal Basis

A lawful basis for processing personal data must be identified and documented before an AI system is developed, tested, or deployed. The following requirements apply:

- If data is genuinely anonymised to ICO standards, UK GDPR does not apply.
- If there is any reasonable risk of re-identification, the data must be treated as personal data, and a lawful basis under Data Protection Act 2018 is required.
- This requirement applies to all stages of AI use, including:
    - Development and testing of AI systems.
    - Training AI models, even when using anonymised datasets (as this constitutes data processing).

3. Data Minimisation and Anonymisation

- Apply data minimisation principles: use only the data necessary for the defined purpose.
- Ensure anonymisation meets ICO standards and regularly assess the risk of re-identification.
- If anonymisation cannot be guaranteed, treat the data as personal and apply appropriate safeguards.

4. Fairness and Transparency

- Be transparent with patients, staff, and stakeholders about how AI is used and what data is processed.
- Staff must not input personal data into public generative AI models (ICO, 2024).
- Ensure AI systems operate fairly, avoiding bias and discrimination.
- Provide clear information on decision-making processes where AI influences outcomes.

- Where appropriate, BrisDoc will undertake Equality Impact Assessments (EqIAs) of AI use-cases, particularly where they may differently affect protected groups.

5. Healthcare Applications of AI

This policy applies to BrisDoc's operational, clinical, and administrative uses of AI. It does not ordinarily cover product development, algorithmic research, or externally led research and development unless BrisDoc acts as the data controller or research sponsor.

For the purposes of this policy, healthcare-related AI use falls into three broad categories:

1. AI specifically designed for healthcare settings (e.g., clinical decision support tools).
2. AI applied to population health or research (e.g., predictive analytics for public health trends).
3. Freely available or commercial 'generic' AI systems (e.g., language models for administrative tasks).

Procurement and Implementation of AI Products

The following requirements must be met when procuring or implementing an AI product or system:

1. Follow the procurement process set out in BrisDoc's Procurement Policy.

2. Engage the Digital, Technical and Business Intelligence Teams to assess architecture, integration, security, interoperability, support, and compatibility with existing infrastructure.

3. Complete and approve a Data Protection Impact Assessment (DPIA) before implementation. The responsible service area and supplier must participate fully in the assessment.

4. Use the ICO's AI and data protection guidance and risk toolkit to identify and reduce data protection risks.

5. Ensure all clinical or patient-facing AI meets the applicable NHS England Digital Technology Assessment Criteria (DTAC) before deployment.

6. Record and address ethical concerns, potential bias, and wider societal impacts through the DPIA, DTAC, and other relevant assurance processes.

7. Obtain the required Health Research Authority approvals for AI used in research projects.

8. Consult the Clinical Safety Officer and, where appropriate, the Medical Device Safety Officer throughout procurement, implementation, operation, change, and decommissioning.

9. Where an AI product is classified as a medical device, confirm that it has the required valid UKCA marking and meets applicable MHRA requirements.

10. Ensure compliance with the UK GDPR requirements governing solely automated decision-making that produces legal or similarly significant effects.

11. Ensure that a suitably qualified person reviews AI-generated outputs before any decision is finalised. BrisDoc will maintain meaningful human oversight for decisions affecting patient care, safety, employment, or other significant outcomes.

Additional Requirements for AI Use

Incident reporting: Any concern, error, near miss, or adverse event involving an AI system must be reported in accordance with BrisDoc's Incident Reporting Policy.

Incident response plans must address foreseeable security, clinical, information governance, and operational events, including data breaches, unsafe outputs, system failures, and loss of service.

Transparency: BrisDoc must communicate relevant AI use clearly to staff and patients, including how the technology is used and how it may affect care, services, or employment. Personal data must not be entered into public generative AI tools.

Ethical use of data: Data processed by AI systems must be handled lawfully, fairly, securely, and transparently, with appropriate consent and anonymisation controls where required. Consent is required where the AI use is optional or falls outside direct care, in line with applicable ICO and NHS England guidance.

Secure data storage: Data used or produced by AI systems must be stored securely and protected throughout its lifecycle through proportionate access controls, retention arrangements, monitoring, and secure disposal.

Patient and public involvement: Appropriate engagement should be undertaken where AI use may materially affect patients or the public, so that concerns about data use, fairness, transparency, and service impact can be understood and addressed.

AI-related incidents with actual or potential patient impact must also be reviewed through the relevant clinical safety processes, including DCB0129 and DCB0160 where applicable.

Consultation for service changes: Where AI implementation changes services, roles, or working practices, BrisDoc must consult affected employees and recognised trade unions in accordance with its organisational change policy.

Bias and discrimination: AI products must be assessed before deployment and monitored in use to identify and mitigate unfair bias, discriminatory outcomes, and unequal effects on protected or underserved groups.

Continuous monitoring: AI systems must be monitored throughout their lifecycle for security events, anomalies, inappropriate use, performance degradation, bias, and other indicators of harm or non-compliance.

Training

A structured training programme will ensure that staff involved in selecting, implementing, using, or overseeing AI understand their responsibilities and can use these technologies safely, responsibly, and in accordance with organisational requirements.

Induction Training

New starters will receive an introduction to the organisation's approach to AI during their induction period. This will include an overview of how AI is used within services, key governance principles, and expectations around safe and compliant use. Induction materials will be updated regularly to reflect changes in policy, technology, or best practice.

Annual Refresher Training All staff must complete annual refresher training covering current guidance, emerging risks, and changes to organisational policy. The training will reinforce regulatory responsibilities and expectations for the appropriate use of AI tools.

Role-Specific Training Training must reflect the responsibilities and risks associated with different roles:

- Clinical staff: safe integration of AI into clinical decision-making, understanding model limitations, escalation routes, and maintaining professional standards.
- Administrative staff: appropriate use of AI for workflow automation, communication support, or efficiency tools, with emphasis on data handling and accuracy.
- Digital, BI, and Innovation teams: deeper technical training covering implementation, evaluation, monitoring, risk management, and ongoing model governance.

Prohibition and Responsible Use Training All staff must be trained on prohibited practices, including the requirement not to enter personal, confidential, or sensitive information into public or non-approved AI systems. Training will reinforce data protection duties, organisational security requirements, and the potential consequences of non-compliance.

Ongoing Communication Regular staff briefings, newsletters, and targeted communications will keep colleagues informed about governance requirements, recognised good practice, emerging risks, and material changes in the AI landscape.

Consultation

Relevant stakeholders must be involved in the development, approval, implementation, management, and monitoring of AI under this policy. This includes the SIRO, DPO, Executive Sponsor, Information Governance Lead, Research and Development Team where research is involved, Procurement Team, Clinical Safety Officer, Business Intelligence, Digital, and Technical teams, together with affected service leads and workforce representatives where appropriate. This collaborative approach supports proportionate assurance and alignment with BrisDoc's operational and clinical needs.

Disciplinary

All staff, regardless of profession, role, or band, have a legal and contractual responsibility to comply with this policy and with the procedures and standards referenced within it.

BrisDoc treats breaches of confidentiality and other failures to comply with information governance requirements seriously. Staff whose actions breach this policy or related policies may be subject to disciplinary action and, where applicable, investigation or prosecution under the Data Protection Act 2018 and UK GDPR.

Where a breach appears intentional, BrisDoc will consider whether notification or referral to the Information Commissioner's Office, the police, or another competent authority is required. Such a breach may result in criminal proceedings against the individual concerned.

Dissemination

This policy will be made available through BrisDoc's approved policy channels and communicated to staff through appropriate internal briefings and management communications. Managers are responsible for ensuring that their teams know where to find the policy, understand its key requirements, and complete mandatory AI governance training.

Process for Monitoring

This policy will be reviewed regularly and updated to reflect changes in legislation, national guidance, technology, and organisational need. Monitoring will provide assurance that AI systems remain safe, lawful, effective, and aligned with BrisDoc's governance standards. The following criteria define how compliance will be assessed.

Monitoring and Audit Criteria

Completion of DPIAs: A DPIA must be completed, reviewed, and approved before any AI system is implemented. Assurance will include review of the DPIA register, project documentation, and proportionate spot checks during implementation and at least annually thereafter. The Information Governance Lead and DPO are responsible for this oversight.

DTAC evidence: Relevant digital health technologies must meet applicable NHS England DTAC requirements before deployment. Supplier evidence will be reviewed and compliance validated during procurement and annual reassessment by the Digital and Technical Teams, Clinical Safety Officer, and Information Governance Lead.

Annual review of AI products: Each AI product in use must undergo an annual review covering safety, accuracy, performance, bias, information governance, and service impact. Reviews will draw on performance data, incidents, user feedback, supplier information, and discussion with service leads. Digital, Business Intelligence, Clinical Safety, Information Governance, and service leadership teams share responsibility for this review.

Clinical safety sign-off: Any AI tool with actual or potential patient impact must have the required clinical safety documentation, including DCB0129 and DCB0160 where applicable, completed and approved before go-live. Clinical safety files and sign-off must be reviewed at implementation and following material changes. The Clinical Safety Officer, Caldicott Guardian, and SIRO are responsible for assurance.

Monitoring of audit logs: Appropriate audit logs must be retained and reviewed to identify inappropriate access, misuse, anomalous behaviour, or potential security events. Monitoring will combine automated analysis with periodic manual review and quarterly reporting. IT Security, Cyber Security, and Digital Teams are responsible for this activity.

Incident themes review: AI-related incidents, near misses, and concerns must be reviewed to identify recurring causes and themes, including clinical safety, bias, data protection, security, and operational resilience. Thematic reviews will take place quarterly and after any significant incident, with learning shared through relevant governance forums. Clinical Safety, Information Governance, and service governance teams are responsible.

Appendix 1 – Terms

Artificial Intelligence (AI) and Generative Artificial Intelligence: AI refers to the development of computer systems capable of performing tasks that typically require human intelligence, such as visual perception, speech recognition, decision-making, and translation. Generative AI is a subset of AI that creates new or unique data by identifying patterns within input data.

Data Protection Impact Assessments (DPIA): A DPIA is used to identify and minimise the data protection risks of a project. BrisDoc mandates DPIAs for any new data processing activities, systems, services, or commissioning. The Information Governance Lead will review, approve, and provide guidance on DPIAs where necessary.

Digital Technology Assessment Criteria (DTAC): DTAC is an NHS-developed framework for assessing digital health technologies. It includes criteria related to clinical safety, data protection, technical security, interoperability, and usability. Any digital health product used within BrisDoc must meet DTAC standards.

Machine Learning: A subfield of AI that enables computers to learn and adapt without explicit instructions, using algorithms and statistical models to identify patterns in data and make decisions. Machine learning can be used for tasks such as image categorisation, data analysis, and predictive modelling.

Natural Language Processing (NLP): A branch of AI focused on enabling computers to understand and process human language, whether written or spoken, in a manner similar to human comprehension.

Processing: Refers to any action performed on data, including obtaining, recording, using, altering, storing, retrieving, disclosing, or erasing the data. Any interaction with data is considered "processing."

Robotic Process Automation (RPA): A technology that automates routine business tasks by emulating human actions, such as data entry and form completion. Unlike AI, RPA is driven by predefined processes. However, AI and RPA can work together, with AI handling data-driven tasks and RPA automating process-driven tasks.

Solely automated decision-making – decisions made without any human involvement that have legal or significant effects on an individual.

Version Control

Date	Version	Author	Change Details
10/02/2026	1.0	DL	Document created.
23/10/2026	2.0	JR	Re-write. New policy.