

Explainable AI Engine Service Definition Document

What the service is:

Our Explainable AI (XAI) service provides comprehensive artificial intelligence solutions specifically designed for healthcare organisations, with built-in transparency and interpretability to support clinical decision-making. The service enables NHS trusts, private healthcare providers, and medical research institutions to deploy machine learning models for diagnostic support, treatment planning, patient risk stratification, and clinical workflow optimisation whilst maintaining full understanding of how AI-driven recommendations are made. This ensures compliance with NHS AI Lab guidelines, NICE evidence standards, MHRA regulations for software as a medical device, and builds trust with clinicians and patients. Healthcare-specific features include integration with HL7 FHIR standards, SNOMED CT terminology support, and DICOM compatibility for medical imaging. The service includes a clinician-friendly dashboard for visualising model predictions, feature importance analysis showing which clinical factors influence decisions, decision boundary exploration, and counterfactual explanations demonstrating how changes in patient parameters would affect outcomes. It supports integration with existing electronic health record (EHR) systems, PACS, and clinical data warehouses through secure REST APIs and can be deployed on-premises within hospital infrastructure or in NHS-approved cloud environments.

• Data Backup, Restore and Disaster Recovery

All healthcare data including AI models, clinical training datasets, patient explanation outputs, and configuration settings are backed up automatically in compliance with NHS Digital Data Security and Protection Standards and the Data Security and Protection Toolkit (DSPT) requirements.

Levels - Tier 1 - Real-time replication with continuous data protection and point-in-time recovery capability down to 5-minute intervals for critical patient safety systems; Tier 2 - Hourly incremental backups with 30-day retention for operational clinical data; Tier 3 - Daily full backups retained for 90 days with monthly archival backups retained for the statutory 8 years minimum (extendable to meet specific NHS records retention schedules). All backups containing patient identifiable data are pseudonymised where appropriate and encrypted at rest using AES-256 encryption with NHS-approved key management, stored in geographically separate UK data centres that meet NHS infrastructure security standards.

Backup integrity is verified through automated monthly restoration tests. Restore operations can be initiated through the secure management console with dual authorisation for patient data, with typical restoration times of under 2 hours for critical clinical systems and under 30 minutes for individual model or dataset recovery. All

backup and restore activities are logged in tamper-proof audit trails accessible for NHS information governance reviews.

Disaster recovery

Our Business Continuity Plan ensures uninterrupted clinical service availability through redundant infrastructure across multiple UK availability zones compliant with NHS Digital Technology Infrastructure standards, automated failover mechanisms with sub-30 second switching times for patient-critical systems, and regular quarterly testing of recovery procedures including simulated major incident scenarios.

The Disaster Recovery Plan maintains a Recovery Time Objective (RTO) of 2 hours for clinical decision support systems and 4 hours for non-critical services, with a Recovery Point Objective (RPO) of 5 minutes for patient data to minimise any potential impact on clinical care.

In the event of a major incident affecting our primary data centre, services automatically fail over to our secondary UK-based site with minimal data loss, ensuring continuous availability of AI explanations for ongoing clinical decisions.

All disaster recovery procedures are independently audited annually and comply with ISO22301 business continuity management standards, NHS Digital operational resilience standards, and the Civil Contingencies Act requirements for essential services. Our incident response team maintains 24/7 availability and coordinates with NHS Digital CareCERT for any cybersecurity incidents affecting healthcare data.

• Onboarding and Offboarding Support

Onboarding Support

Our comprehensive healthcare onboarding programme includes a dedicated Clinical AI Implementation Manager and Information Governance Specialist assigned for the first 90 days, with experience in NHS digital transformation projects.

We provide structured training sessions covering platform fundamentals, clinical explainability techniques, NICE AI evidence standards, and best practices for AI interpretation in clinical contexts. Training is tailored for different healthcare roles including clinicians, radiologists, pathologists, clinical informaticians, and IT teams. We provide customised workshops aligned to your organisation's specific clinical use cases (diagnostics, treatment planning, patient risk stratification), technical documentation including HL7 FHIR integration guides and EHR connection examples, and access to our healthcare-specific knowledge base with clinical case studies and video tutorials.

The onboarding process includes an initial clinical safety assessment and information governance review, environment setup within NHS N3/HSCN networks or approved



cloud infrastructure, secure integration with existing EPR/EHR systems and PACS, data protection impact assessment (DPIA) support, user training sessions for clinical,

technical, and administrative staff including safe AI usage protocols, and ongoing clinical validation support during the first quarter. We offer optional advanced training on clinical explainability methods such as counterfactual explanations for treatment decisions, model debugging techniques for clinical AI, MHRA compliance for software as a medical device, and NHS AI Lab assessment framework preparation. All training includes awareness of clinical risk management and the duty of candour when using AI systems.

Offboarding Support

When customers choose to transition away from our service, we provide structured offboarding support to ensure smooth continuity. This includes a 30-day notice period with continued full-service access, comprehensive data extraction in standard formats (CSV, JSON, ONNX for models), detailed documentation of all customisations and configurations, and assistance with data migration to alternative platforms.

We provide export of all historical explanation results, audit logs, and model performance metrics. Our team offers consultation on transitioning to alternative solutions and can provide recommendations for maintaining explainability practices.

All customer data is securely deleted from our systems within 90 days of contract termination, with certified proof of deletion provided upon request. We maintain professional relationships and offer re-onboarding at preferential rates should customers wish to return.

• Service Constraints

Scheduled maintenance windows occur on the second Sunday of each month between 02:00-06:00 GMT, during which the platform may experience brief interruptions. Emergency maintenance is conducted only when critical security patches are required, with a minimum 48-hour advance notification.

The service supports extensive customisation including custom explainability algorithms, branded user interfaces, integration with proprietary ML frameworks, custom report templates, and API endpoint modifications. However, certain core security and compliance features cannot be modified to maintain system integrity and regulatory compliance. Customisations that significantly impact system performance or security posture require approval from our technical review board.

The platform supports multi-tenancy with logical data separation but does not offer dedicated infrastructure for individual customers in the standard service tier.

- **Service Levels**

Performance

The platform delivers API response times under 200ms for 95% of requests and under 500ms for 99% of requests. Explanation generation for standard models completes within 2 seconds for SHAP and LIME methods. The system scales horizontally to

support up to 10,000 concurrent explanation requests with auto-scaling based on demand. Batch processing capabilities handle up to 1 million predictions per hour with explanation outputs.

Availability

We guarantee 99.9% uptime SLA (excluding planned maintenance windows), which equates to a maximum of 43 minutes of unplanned downtime per month. The service is monitored 24/7 with real-time alerting. Service credits are provided for any downtime exceeding the SLA commitment.

Support Hours

Standard support is available Monday to Friday 08:00-18:00 GMT with email and ticketing system access. Priority support (available as an add-on) provides 24/7 coverage including weekends and bank holidays with telephone support and guaranteed response times of 1 hour for critical issues, 4 hours for high priority, and 8 hours for normal priority.

- **After Sales Support**

Ongoing support includes regular software updates and security patches deployed automatically, quarterly business reviews with account managers to assess service performance and identify optimisation opportunities, access to a dedicated support portal with comprehensive documentation and community forums, and proactive monitoring alerts for anomalies in model behaviour or explanation quality.

We provide version upgrade assistance with backward compatibility guarantees for major releases, technical consultation on implementing new explainability techniques, and assistance with regulatory compliance documentation.

Customers receive monthly service reports detailing uptime, performance metrics, and usage analytics. Advanced support tiers include dedicated technical account managers and priority feature requests.

- **Technical Requirements**

The service supports clinical AI models trained in ML frameworks with specific support for medical imaging models (chest X-ray, CT, MRI analysis) and clinical prediction models.

For cloud deployment, the service integrates with NHS-approved cloud platforms including AWS (N3/HSCN connected), Azure (including Azure for Health), and Google Cloud Healthcare API.

On-premises deployment within hospital infrastructure requires Linux-based systems with Docker support, minimum 32GB RAM (64GB recommended for medical imaging workloads), modern CPU with AVX2 instruction set support, and GPU acceleration required for radiology AI models.

Network connectivity requires secure access via NHS N3/HSCN networks or approved VPN solutions, with HTTPS access to our API endpoints using TLS 1.3.

All network communications support NHS Digital security standards including IPsec encryption. For clinical dashboard access, users need NHS-approved modern web browsers (Chrome 90+, Firefox 88+, Edge 90+).

API clients must support REST protocols with HL7 FHIR resource formatting and JSON data structures. Single Sign-On (SSO) integration is available via NHS Identity, SAML 2.0, OAuth 2.0, and OpenID Connect, with support for NHS Care Identity Service authentication.

• Outage and Maintenance Management

Planned maintenance is scheduled during the monthly maintenance window with 7 days advance notice provided via email and the service status page. Emergency maintenance for critical security vulnerabilities is performed with a minimum 48-hour notice period whenever possible. During maintenance, the system may operate in read-only mode or experience reduced capacity rather than complete outage.

In the event of an unplanned outage, our incident response team is automatically alerted within 60 seconds via our monitoring systems. Status updates are published to our status page every 30 minutes during incidents, with email and SMS notifications sent to registered contacts.

Post-incident reviews are conducted for all service disruptions exceeding 15 minutes, with root cause analysis and remediation plans shared with affected customers.

Service credits are automatically applied for any breach of our SLA commitments.

• Hosting Options and Locations

The service is available in three NHS-compliant deployment models:

1. Cloud SaaS hosted on NHS-approved cloud infrastructure with N3/HSCN connectivity in UK regions only (i.e. London, Manchester);
2. Private Cloud deployment within NHS trust's chosen approved cloud provider (AWS with NHS connectivity, Azure for Health, or GCP Healthcare API) with management handled by our clinical informatics team;

3. and On-Premises deployment within hospital data centres with full control over data residency and air-gapped options for highly sensitive research data.

All locations hold ISO27001, ISO27799 (health informatics security), and Cyber Essentials Plus certification, and comply with UK GDPR, Data Protection Act 2018, NHS Digital Data Security and Protection Standards and achieve NHS Data Security and Protection Toolkit (DSPT) Standards Met status.

Healthcare customers can specify their preferred region for patient data residency with contractual guarantees that identifiable patient data never leaves UK jurisdiction.

The service can be deployed behind NHS firewalls with air-gapped deployment available for research environments handling highly sensitive genomic or mental health data.

• Access to Data (upon exit)

Upon contract termination or service exit, NHS trusts and healthcare organisations retain full ownership and control of all patient data, clinical models, and explanation outputs in accordance with NHS records management and data controller responsibilities.

We provide up to 90 days of continued read-only access (extendable to 12 months for clinical continuity if required) to facilitate orderly transition without disrupting patient care. All data including trained clinical models, anonymised training datasets, patient explanation outputs, clinical validation reports, configuration files, and comprehensive audit logs are made available for download in healthcare-standard formats including HL7 FHIR JSON, CSV, ONNX, PMML, DICOM (for imaging models), and HDF5. For patient-identifiable data, exports are provided via NHS-approved secure channels including N3/HSCN SFTP, encrypted cloud storage bucket transfer with NHS-grade encryption, or encrypted physical media delivered by secure courier for large datasets.

We offer clinical transition assistance including documentation mapping to alternative platforms, format conversion services preserving clinical metadata, HL7 FHIR transformation support, and technical consultation on maintaining explainability capabilities and clinical safety in new environments.

Clinical model validation documentation and safety case reports are provided to support continued use or re-validation in new systems.

All patient data is permanently and verifiably deleted from our systems within 90 days of service termination (or as specified by NHS records retention requirements), with certified proof of deletion compliant with NHS Digital standards and Caldicott principles. Healthcare organisations can request earlier data deletion at any time during the retention period following appropriate clinical governance approvals.

For research data, we support transfer to NHS-approved trusted research environments. The offboarding process includes information governance sign-off,

clinical safety confirmation, and handover documentation for incoming system suppliers to ensure continuity of care.

• Security

Our healthcare security framework is certified to ISO27001 (information security), ISO27017 (cloud security), ISO27018 (cloud privacy), ISO27799 (health informatics security), and SOC 2 Type II standards. We achieve NHS Data Security and Protection Toolkit (DSPT) Standards Met status annually and comply with NHS Digital Data Security and Protection Standards including the Data Security and Protection Requirements. All patient-identifiable data is encrypted at rest using AES-256 encryption with NHS-approved key management (stored in UK-based Hardware Security Modules), and in transit using TLS 1.3 with perfect forward secrecy.

We implement multi-factor authentication (MFA) as mandatory for all clinical users, with support for NHS smartcards, hardware security keys, biometric authentication, and FIDO2 standards. Role-based access control (RBAC) aligned to NHS workforce roles provides granular permissions management with separation of duties for clinical, administrative, and technical staff, supporting custom roles and attribute-based access control (ABAC) for complex clinical scenarios.

All access to patient data and system activities are logged with tamper-proof audit trails retained for 8 years minimum (extendable to meet NHS records retention schedules), with real-time alerts for suspicious access patterns.

We conduct annual penetration testing by independent CHECK-certified and CREST-certified security firms with healthcare specialisation, quarterly vulnerability assessments, continuous security monitoring through our NHS-aligned Security Operations Centre (SOC), and regular submissions to the NHS Digital Data Security Centre. Our SOC provides 24/7 threat detection and incident response with direct connection to NHS Digital CareCERT for healthcare-specific cyber threats.

Patient data isolation between organisations is maintained through cryptographic separation at the application layer with additional physical segregation options for highly sensitive data. We maintain Cyber Essentials Plus certification, comply with the UK GDPR, Data Protection Act 2018, the Network and Information Systems (NIS) Regulations, the Health and Social Care Act information governance requirements, and the common law duty of confidentiality.

Security patches for patient-facing systems are applied within 24 hours of release for critical vulnerabilities affecting clinical safety or data protection. All personnel undergo enhanced DBS checks, NHS-approved information governance training annually, and clinical staff receive additional training on patient confidentiality and Caldicott principles. We maintain comprehensive clinical risk management documentation (DCB0129 and DCB0160 compliant), conduct regular Data Protection Impact Assessments (DPIAs) for new features, and support NHS trust information governance frameworks including Caldicott Guardian oversight.