Sword Intelligence Service Definition

What is Sword Intelligence

Sword Intelligence is an AI transformation partner for NHS trusts and healthcare organisations, delivering composable AI agents that automate administrative workflows, patient engagement, and care coordination. Our platform acts as an intelligent layer that integrates with your existing systems to handle routine tasks, allowing your staff to focus on complex cases and direct patient care.

What makes our approach different is flexibility. Rather than replacing your systems or forcing you to adapt to rigid software, our agents integrate into your current environment—whether that's your EPR, PAS, contact centre infrastructure, or patient-facing applications. The platform supports multiple communication channels including voice, SMS, email, and web-based interactions, ensuring patients can engage in ways that work for them.

Data Backup, Restore, and Disaster Recovery

We maintain robust backup and disaster recovery processes designed to protect your data and ensure business continuity. Our infrastructure includes automated daily backups with point-in-time recovery capabilities, geographically distributed data storage for resilience, and documented disaster recovery procedures with defined recovery time objectives (RTOs) and recovery point objectives (RPOs).

Specific backup schedules, retention periods, and disaster recovery targets are defined collaboratively during implementation to align with your organisation's data governance requirements and operational criticality. For live clinical deployments, we work with you to ensure recovery processes meet your service continuity obligations.

All backup and recovery processes are tested regularly, and we can provide documentation of these tests upon request as part of your assurance requirements.

Onboarding and Offboarding Support

Onboarding: Every deployment begins with a structured onboarding process tailored to your scope and complexity. All customers receive documentation and resources appropriate to their deployment, which may include configuration guides, technical documentation, operational materials, and training content for both technical and clinical/operational users.

For live services or larger deployments, we provide guided onboarding sessions delivered remotely to support initial setup, configuration, and testing. These sessions typically involve technical leads, service owners, and relevant clinical or operational stakeholders. Where needed, we can also provide tailored training for administrators and frontline users, focused on day-to-day operations, governance, and change management. For complex or system-wide implementations, onsite support and in-person training can be arranged.

You're assigned a dedicated Technical Account Manager from day one who serves as your primary point of contact throughout deployment and ongoing operations. Ongoing support continues throughout rollout to ensure teams are confident using the service and can adopt it safely and effectively.

Offboarding: End-of-contract processes, including data extraction, transition support, and knowledge transfer, are discussed and agreed during contract negotiations. We work collaboratively to ensure smooth transitions that align with your technical environment, data governance requirements, and operational needs. Data extraction methods, formats, and any associated support are defined on a per-contract basis to ensure you maintain full control of your information.

Service Constraints and Customisation

Maintenance Windows: Planned maintenance is limited to a maximum of 4 hours per month and scheduled during off-peak hours whenever possible. We provide at least 48 hours advance notice via email before any planned maintenance, and use commercially reasonable efforts to minimize service disruption during these windows.

Customisation: Our AI agents are designed to be fully tailored to your workflows, standard operating procedures, and local policies. Almost everything can be customized, including triage pathways and routing rules, conversation flows and messaging, operating hours and escalation logic, integration points with your existing systems, and reporting outputs. The platform adapts to your local services, referral options, and operational requirements.

Straightforward configurations—such as operating hours, user permissions, and basic workflow rules—can be adjusted directly through the platform interface by authorized users. More sophisticated customizations, such as complex routing logic, custom integrations, or clinical pathway modifications, are implemented collaboratively with our support engineering team to ensure safety and quality.

All changes are version-controlled, tested in non-production environments before deployment, and can be rolled back if needed. Role-based access controls ensure only authorized users can make changes, maintaining governance and accountability while giving you flexibility to evolve the service as your needs change.

Service Levels: Performance, Availability, and Support Hours

Platform Availability: We target 99.5% uptime, calculated excluding planned maintenance windows and circumstances outside our direct control (such as client infrastructure issues, third-party service failures, or force majeure events). Uptime is measured yearly and forms the basis for any service credit eligibility.

Support Hours and Response Times: All clients receive comprehensive support that scales with deployment needs. Standard support includes email and ticketing support during business hours (Monday-Friday, 9am-5pm GMT). Response times are based on severity:

- **Critical incidents** (platform down or unavailable): 2-hour response target, 24/7 support
- **High severity** (significant performance impairment): 4-hour response target, 24/7 support
- **Medium severity** (partial, non-critical loss of functionality): 1 business day response
- **Low severity** (minor issues): 3 business days response
- **Suggestions for improvement**: 5 business days response

For Critical and High severity incidents, we provide 24/7 support year-round, ensuring urgent production issues are addressed immediately regardless of time or day. Your dedicated Technical Account Manager will reach out proactively for severe incidents to keep you informed throughout resolution.

Resolution targets for Critical incidents are 24 hours, and High severity incidents are 48 hours, measured from when we receive sufficient information to investigate. For Medium and Low severity issues, we assess circumstances and communicate resolution timelines directly.

Performance: Platform performance metrics—such as call handling times, API response latencies, and workflow processing speeds—are monitored continuously and reported based on your deployment requirements. Expected performance levels are defined during implementation planning based on your anticipated volumes and use cases.

After-Sales Support

After deployment, you continue to work with your dedicated Technical Account Manager and have access to our support engineering team for technical escalations, infrastructure questions, and performance optimization.

We provide ongoing access to documentation updates, product enhancement notifications, and opportunities to provide feedback on roadmap priorities. For customers operating at scale or with system-wide deployments, we can arrange regular service reviews to discuss performance, emerging needs, and strategic alignment.

Enhancement requests and suggestions for improvement are reviewed and considered for inclusion in our product roadmap. We acknowledge receipt and provide feedback on feasibility and potential timelines, though we're under no obligation to implement suggested enhancements.

Technical Requirements

Our flexible, composable architecture is built to work with any existing system, so there are no

specific technical prerequisites. The platform is browser-based and works with all major browsers (Chrome, Firefox, Safari, Edge) on any operating system.

Depending on deployment requirements, we can integrate directly into your existing EPR or PAS systems, allowing users to access functionality within their current workflows without switching applications. Integration is achieved through APIs that connect with NHS systems including EPRs, PAS platforms, and third-party applications.

For voice-based deployments, the service works through standard telephony infrastructure or can integrate with your existing contact centre systems. SMS and email channels work through standard protocols. Where you have existing patient-facing mobile applications, we can integrate our AI agents directly into those platforms.

API specifications and integration requirements are defined collaboratively during deployment planning to ensure seamless integration with your specific technical environment.

Outage and Maintenance Management

Unplanned Outages: Service outages and degradations are communicated proactively through email notifications to affected users. For Critical or High severity incidents, your Technical Account Manager also reaches out directly to ensure you're informed and updated throughout resolution.

Where required for your incident management processes, we can provide status updates via API or through a status dashboard. Communication methods are confirmed during implementation to align with your operational needs.

All outages are investigated, and post-incident reports are produced for significant incidents, summarizing what occurred and remediation steps taken.

Hosting Options and Locations

Sword Intelligence operates as a public cloud service, hosted on secure, enterprise-grade infrastructure. Data residency and hosting locations are discussed during contract negotiations to align with your data governance, regulatory, and information security requirements.

For UK NHS deployments, we can ensure data remains within UK boundaries and meets NHS Data Security and Protection Toolkit requirements. Specific hosting arrangements, including any multi-region redundancy or specific geographic constraints, are agreed on a per-contract basis.

Infrastructure security includes encryption at rest and in transit, network isolation, regular security assessments, and compliance with recognized standards including HITRUST CSF, SOC 2, and ISO/IEC 27001.

Access to Data Upon Exit

Data extraction processes and formats are discussed and agreed during contract negotiations to ensure they align with your technical environment, data governance requirements, and operational needs. We work collaboratively to ensure you maintain full control of your information and can transition smoothly to alternative solutions or internal systems.

Security

Information Security Programme: Sword Health's security programme is built on the HITRUST Common Security Framework (CSF) and reviewed annually through HITRUST and SOC 2 audits. Our policies cover data classification, incident management, secure development lifecycle (SDLC), vulnerability management, and access control. The Chief Information Security Officer (CISO) approves and oversees these policies and leads the security and incident response teams.

Configuration and Change Management: All configuration and software changes are managed using Infrastructure as Code and Configuration as Code. All code changes are tracked using Jira tickets and subject to mandatory code reviews, testing, and approval by a team leader before production deployment. Code security tools are embedded in the CI/CD pipeline.

Vulnerability Management: We classify vulnerabilities by severity and commit to fixing critical findings within 7 days, high severity within 30 days, and medium severity within 90 days. Threats are identified through annual penetration tests, static code analysis (Wiz for cloud infrastructure, Snyk for application code), continuous monitoring tools (AWS GuardDuty, Google Cloud Security Command Center), and external scans.

Protective Monitoring: We use a managed detection and response service (Expel) that ingests logs from devices, cloud platforms, identity systems, and other tools to identify indicators of compromise. The Computer Security Incident Response Team (CSIRT) investigates incidents, isolates affected systems, disables compromised accounts, and applies fixes. Response targets are tied to severity: critical incidents within 2 hours, moderate within 24 hours, minor within 3 days.

Incident Management: Our incident management plan defines playbooks for typical events (ransomware, stolen credentials, data loss, API compromises) with assigned severities determining response windows. Users report suspected incidents to the Information Security team within 6 hours of discovery. Each incident is logged and, once closed, a post-incident report summarizes what happened and remediation steps taken.

Authentication and Access Control: User authentication requires Multi-Factor Authentication (MFA) and username/password credentials. Access to management interfaces and support channels is restricted using strong authentication and role-based access control, with

least-privilege permissions, audit logging, and secure administrative access mechanisms.

Audit and Logging: Buyers control when and how they access audit information about user actions and supplier actions. Audit data and system log retention periods are user-defined based on your organizational requirements and regulatory obligations.

Compliance Standards: Our security practices comply with recognized standards including:

- HITRUST CSF
- SOC 2 Type II
- ISO/IEC 27001
- CSA CCM v4.0
- NHS Data Security and Protection Toolkit