converse360

Information Security Policy

Assist-Me AI Agent Platform

Amendment History

Ver	Date	Comments	Modified by
1.0	3 August 2024	Initial release	Dave Jani
1.1	18 August 2025	Updated Generative AI information	Paul Williams
1.2	26 March 2026	Review and reformatted	Dave Jani

1. Introduction and Purpose

Converse360 is a UK-based Software-as-a-Service (SaaS) provider of AI Agent technology. We develop, operate and support AI-powered automation tools used by our customers to manage business workflows, handle data, and interact with third-party systems.

This Information Security Policy sets out the Company's commitment to protecting the confidentiality, integrity and availability of all information assets, including customer data, AI model configurations, system infrastructure, and proprietary algorithms - from unauthorised access, misuse, disclosure, disruption, modification or destruction.

This Policy applies across the full lifecycle of information at the Company: from collection and processing through to storage, transmission and deletion. It forms the foundation of the Company's Information Security Management System (ISMS) and is aligned with:

- UK General Data Protection Regulation (UK GDPR) and the Data Protection Act 2018
- ISO/IEC 27001:2022 — Information Security Management Systems
- NCSC Cyber Essentials and Cyber Essentials Plus
- NCSC Cloud Security Principles
- The Network and Information Systems (NIS) Regulations 2018 (where applicable)
- AI-specific guidance from the UK ICO and the Alan Turing Institute's AI ethics principles

2. Scope

This Policy applies to:

- All employees, directors, and contractors of converse360
- All third-party suppliers, sub-processors, and technology partners who access, store or process Company or customer data
- All information assets owned, leased or managed by the Company, including cloud infrastructure, SaaS platform components, AI models, APIs, and customer data environments
- All geographic locations from which Company operations are conducted

For the avoidance of doubt, this Policy covers all data processed by the AI Agent platform, including:

- Customer-submitted data (inputs to AI Agents)
- AI model outputs and decisions
- Training data and model weights
- Integration credentials and API keys for third-party systems
- Platform telemetry and usage data

3. Roles and Responsibilities

Role	Responsibility
Board of Directors / MD	Ultimate accountability for information security. Approve this Policy and ensure adequate resources are allocated.
Technical Director / Security Lead	Owns this Policy. Maintain the ISMS. Report security posture to the Board. Lead incident response. Ensure compliance with UK GDPR. Advise on data protection impact assessments (DPIAs). Act as ICO contact.
AI / ML Teams	Ensure AI model training, deployment and inference operations are conducted securely and ethically. Document model data lineage.
Engineering & DevOps Teams	Implement security controls in the platform. Conduct code reviews. Manage cloud infrastructure securely. Follow secure development lifecycle (SDL).
Customer Success / Sales	Ensure customers are informed of security capabilities and obligations. Do not commit to security controls outside agreed scope.
All Staff	Complete security awareness training. Report incidents immediately. Follow this Policy and all related procedures.

4. Information Classification

All information assets must be classified according to the following scheme and handled accordingly:

Classification	Description	Examples	Controls Required
PUBLIC	Information approved for external release.	Marketing materials, public API documentation	Standard access controls
INTERNAL	General business information not for external distribution.	Internal procedures, staff lists, meeting notes	Access limited to staff and authorised contractors
CONFIDENTIAL	Sensitive business or customer information. Disclosure would cause significant harm.	Customer data, contracts, financial records, AI model configs	Encrypted at rest and in transit. Need-to-know access only. Logged access.
RESTRICTED	Highly sensitive. Disclosure would cause severe harm or regulatory breach.	Credentials, API keys, model weights, personal data, security configurations	Strict access control, MFA enforced, encrypted storage, no external transfer without DPA.

NOTE: *All customer data processed by the AI Agent platform is classified as CONFIDENTIAL as a minimum, and RESTRICTED where it constitutes personal data under UK GDPR.*

5. Access Control

5.1 Principles

- Least Privilege: Users are granted the minimum access required to perform their role.
- Need to Know: Access to sensitive information is restricted to those with a documented business need.
- Separation of Duties: No single individual should have end-to-end control of critical processes without oversight.
- Zero Trust: All access requests are verified regardless of network location.

5.2 Authentication

- Multi-factor authentication (MFA) is mandatory for all staff accessing Company systems, cloud infrastructure, and production environments.
- Passwords must meet the NCSC password guidance: a minimum of 12 characters, use of passphrases encouraged, no mandatory periodic rotation unless compromise is suspected.
- Shared accounts are prohibited. Each user must have a unique, personal account.
- Single Sign-On (SSO) must be used where available and supported.

5.3 Privileged Access

- Privileged access (administrative rights to cloud infrastructure, databases, AI infrastructure) must be formally approved.
- Just-in-time (JIT) access provisioning must be used for production environment access where technically feasible.
- Privileged access sessions must be recorded where practicable.

5.4 Joiners, Movers, Leavers

- Access rights must be provisioned within 1 business day of role commencement and removed within 1 day of departure.
- Access rights must be reviewed quarterly by line managers and immediately upon role change.
- A formal leaver process checklist must be completed for all departing staff.

6. Data Protection and UK GDPR Compliance

6.1 Data Protection Principles

The Company processes personal data in accordance with the six principles set out in Article 5 of UK GDPR:

1. Lawfulness, fairness and transparency
2. Purpose limitation
3. Data minimisation
4. Accuracy
5. Storage limitation
6. Integrity and confidentiality

6.2 Lawful Basis for Processing

The Company must document the lawful basis for all personal data processing activities in its Record of Processing Activities (RoPA). For AI Agent operations, the primary lawful bases are:

- Contractual necessity - processing customer data to deliver the contracted SaaS service
- Legitimate interests - platform security monitoring, fraud prevention, and service improvement
- Legal obligation - compliance with regulatory requirements

Where processing relies on consent, that consent must be freely given, specific, informed and unambiguous.

6.3 AI-Specific Data Protection Requirements

Given the nature of the Company's AI Agent platform, the following additional requirements apply:

- A Data Protection Impact Assessment (DPIA) must be completed before deploying any new AI model or AI Agent capability that processes personal data.
- AI model training must not use customer personal data without explicit contractual permission and a documented legal basis.
- Where AI models make or influence decisions about individuals (e.g., automated responses, recommendations), customers must be informed of this and their rights under Article 22 UK GDPR must be respected.
- Data used to train or fine-tune AI models must be documented and it must not be retained beyond its stated purpose.
- Model outputs must not be used to infer special category data without explicit consent and a Schedule 1 condition under the DPA 2018.

6.4 Data Retention and Deletion

- Customer data will be retained only for the duration of the contract plus 3 months for legal/dispute purposes, unless a longer retention period is required by law.
- Upon contract termination, customer data must be securely deleted or returned within 30 days of written request.
- AI model interactions and logs containing personal data must be subject to defined retention schedules.
- Secure deletion procedures must be applied to all RESTRICTED and CONFIDENTIAL data.

7. Cryptography and Encryption

- All data classified as CONFIDENTIAL or RESTRICTED must be encrypted at rest using AES-256 or equivalent.
- All data in transit must be protected using TLS 1.2 or higher. TLS 1.0 and 1.1 are prohibited.
- API keys, secrets, and credentials must be stored in an approved secrets management tool and must never appear in source code, logs or configuration files in plaintext.

8. Network and Cloud Infrastructure Security

8.1 Cloud Security

The Company's AI Agent platform is operated on cloud infrastructure. The following controls apply:

- Cloud deployments must be configured in accordance with the cloud provider's security best practices and the NCSC Cloud Security Principles.
- Public access to cloud storage buckets, databases, and administrative interfaces is prohibited. All such resources must be private by default.

- Production, staging, and development environments must be strictly separated with no shared credentials or access paths.

8.2 Network Controls

- Network segmentation must be implemented to isolate the AI platform, customer data environments, and internal corporate systems.
- Each customer's instance will be held within a separate subnet in AWS. No structured personal data is stored (only conversational), individual DB instances are available upon request.
- Firewall rules must follow a default-deny posture. All permitted rules must be documented, approved, and reviewed quarterly.
- VPN or equivalent zero-trust network access must be required for all remote administrative access.

8.3 API Security

Given the API-centric nature of the AI Agent platform, the following controls are mandatory:

- All APIs must implement authentication (OAuth 2.0 / API key with appropriate scoping) and authorisation.
- API rate limiting and throttling must be implemented to prevent abuse and denial-of-service.
- API inputs must be validated and sanitised to prevent injection attacks.
- All API activity must be logged and monitored for anomalous behaviour.

9. Secure Development Lifecycle (SDL)

Security must be integrated throughout the software and AI development lifecycle - not treated as a final step before release.

9.1 Development Standards

- All developers must complete secure coding training relevant to their technology stack on joining and annually thereafter.
- The Company's development must follow OWASP Secure Coding Practices as a baseline.
- Peer code review is mandatory for all changes to production codebases. Security considerations must be explicitly addressed in review.
- Secrets and credentials must never be committed to source code repositories. Pre-commit hooks and repository scanning must be implemented to detect and prevent this.

9.2 AI Model Security

The Company acknowledges the specific security risks associated with AI systems, including prompt injection, model poisoning, data leakage via model outputs, and adversarial inputs. The following controls apply:

- AI models must be tested for prompt injection vulnerabilities before deployment and after significant updates.
- Model training pipelines must be treated as production systems with equivalent security controls.
- Model outputs must be filtered and validated before being returned to end users or passed to downstream systems.
- The Company must maintain a log of all AI models in production, including version, training data provenance, and known limitations.
- AI models must not be granted direct access to production systems or data stores without human-in-the-loop approval controls or strict guardrails.

10. Third-Party and Supply Chain Security

The Company uses third-party cloud providers, AI model APIs (e.g. LLM providers), software components, and professional services. All third parties who access, process or store Company or customer data must meet the Company's security standards.

- All sub-processors processing personal data on the Company's behalf must be trusted and selected based on security, reliability and compliance standards.
- New third parties must be assessed proportionately to the sensitivity of data accessed and the criticality of the service.
- Third-party AI model and LLM API providers must be assessed for their data handling practices, particularly with respect to whether customer prompt data is used for model training. Contractual protections must be in place.
- A register of all third-party processors and sub-processors must be maintained and published to customers upon request.

11. Security Incident Management

11.1 Incident Definition

A security incident is any event that has, or could have, an adverse effect on the confidentiality, integrity or availability of information assets. This includes, but is not limited to: data breaches, ransomware attacks, unauthorised access, AI model compromise, denial of service, and loss of devices.

11.2 Reporting

- All staff must report suspected security incidents immediately to the Operations Director
- No member of staff should attempt to investigate or remediate a suspected incident independently without the Operations Directors authorisation.

11.3 UK GDPR Breach Notification

- Personal data breaches must be reported to the Information Commissioner's Office (ICO) within 72 hours of becoming aware where the breach is likely to result in a risk to individuals' rights and freedoms.
- Where the breach poses a high risk to individuals, those individuals must be notified without undue delay.
- All personal data breaches, whether notified to the ICO or not, must be recorded in the Company's breach register.

12. Business Continuity and Disaster Recovery

- Systems are designed for resilience and high availability, backup and recovery processes are in place where applicable
- All customer data must be backed up at minimum daily. Backup integrity must be tested at minimum quarterly.
- The AI Agent platform must be designed with high availability architecture (e.g., multi-availability zone deployment, auto-scaling, load balancing) to minimise single points of failure.

13. Physical and Environmental Security

- As a primarily cloud-based SaaS business, the Company's primary physical security obligations relate to employee devices and office premises.
- All staff must use company-managed devices with full-disk encryption.
- Unattended devices must be locked. Clear desk and clear screen policies apply in all Company offices and when working remotely.
- Company cloud infrastructure is hosted in certified AWS data centres operated by approved cloud providers. Physical security of these facilities is the responsibility of the provider and must be confirmed via their audit reports (SOC 2 Type II, ISO 27001 certificate).

14. Human Resources Security

- Pre-employment screening, including identity verification, right to work checks, and where appropriate DBS checks, must be completed for all staff and contractors prior to commencement.
- Security awareness training is mandatory for all staff on joining and at least annually.
- Disciplinary procedures for deliberate or negligent security breaches are defined in the Company's HR policies.
- Confidentiality agreements must be in place for all employees and contractors with access to CONFIDENTIAL or RESTRICTED information.

15. Compliance, Audit and Review

16.1 Regulatory Compliance

The Company must maintain ongoing compliance with:

- UK GDPR and the Data Protection Act 2018
- ISO/IEC 27001:2022
- NCSC Cyber Essentials
- AI regulation guidance as it develops, including the ICO's guidance on AI and data protection

16.2 Internal Audit

- Internal audits of information security controls must be conducted at least annually.
- Audit findings must be reported to the Board and tracked to resolution.

16.3 Policy Review

- This Policy must be reviewed and updated periodically, or following: a significant security incident, a material change to the business or technology stack, changes to applicable law or regulation, or the introduction of new AI capabilities.
- Policy changes must be approved by the Operations Director and Board/CEO and communicated to all staff.