Simpplr Generative AI Security & Compliance Overview

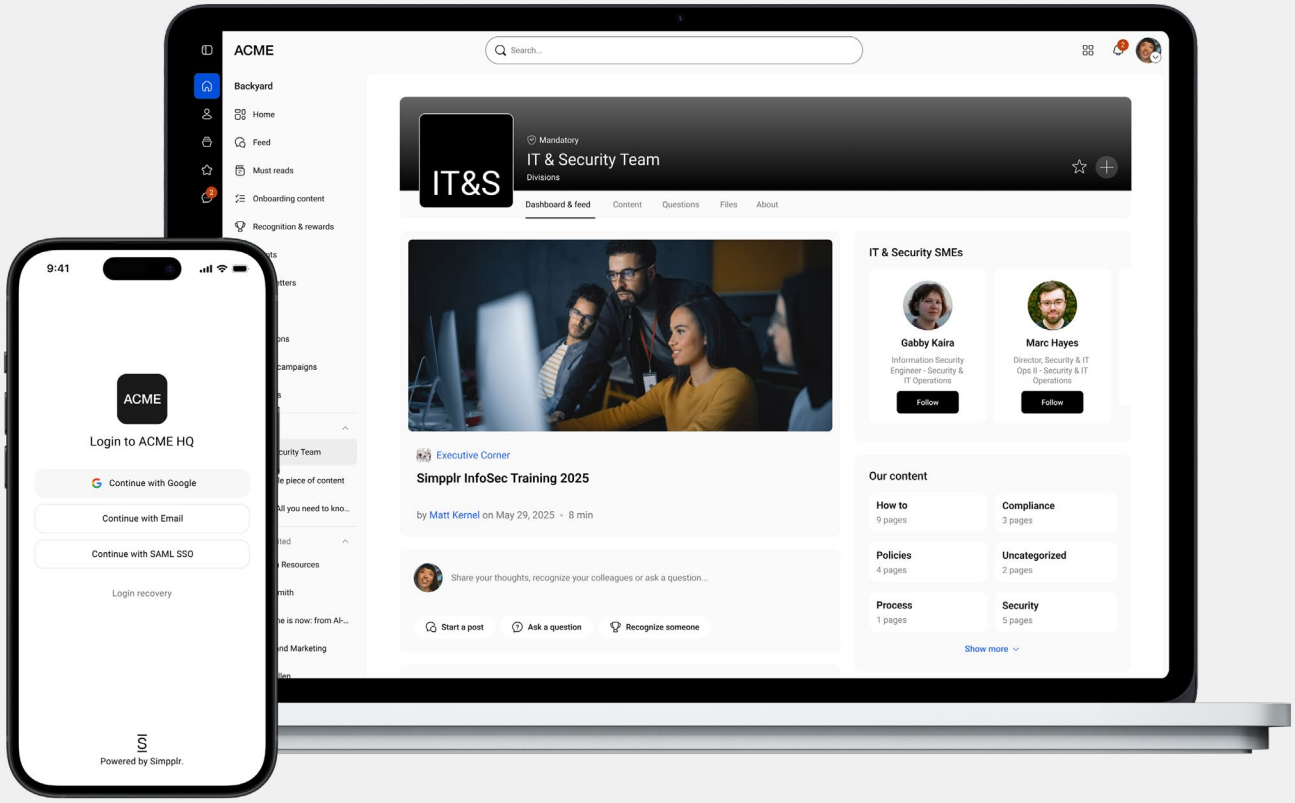

Building trusted, responsible, and secure AI for modern organizations

At Simpplr, we are committed to ensuring that all features, including those powered by generative AI, are developed and operated in a manner that is safe, secure, and compliant.

We believe AI should be both innovative and trustworthy. Simpplr's AI features operate within a security-first, privacy-by-design framework that safeguards customer data, aligns with global standards such as the GDPR, CPRA, and the EU AI Act, and supports enterprise-grade governance and compliance needs. This overview highlights Simpplr's current practices and its ongoing alignment roadmap across AI, privacy, and cybersecurity domains. It describes how these capabilities are implemented securely through defined data-protection measures, regulatory alignment, and governance practices that meet modern enterprise-security expectations.

Best practices for generative AI security

How Simpplr implements AI capabilities securely

Simpplr's generative AI features follow the same disciplined controls that govern its broader information-security framework. Each practice ensures that AI operates safely, responsibly, and in compliance with global standards.

Partner selection

- **Trusted partners:** Simpplr works with partners who commit contractually not to use identifiable customer data.
- **Data isolation:** Customer data is segregated and not mixed with datasets influencing other customers.

Data privacy and confidentiality

- **Usage restrictions:** Customer-confidential data is not incorporated into models unless explicitly approved.
- **Customer consent:** Customers must approve in writing specific AI functionalities; any changes are communicated with advance notice.
- **Third-party subprocessors:** Simpplr uses subprocessors to provide some AI services and discloses them separately, including details of their compliance, security, and certification obligations.

Data retention and deletion

- **Retention:** Data is retained only as long as needed for compliance and operations.
- **Opt-out and deletion:** Customers may opt out of retention and request deletion of their data at any time.

Security measures

- **Access control:** Comprehensive logging and role-based access controls govern AI operations.
- **Private instances:** AI workloads run in private cloud instances within Simpplr's secure infrastructure.
- **Encryption:** All data is encrypted in transit and at rest.
- **Ongoing enhancements:** Simpplr continues to strengthen its AI security through continuous **AI red teaming** and **watermarking** initiatives.

Monitoring and auditability

- Simpplr regularly reviews AI outputs to ensure accuracy and detect undesirable behavior such as hallucinations or bias.
- All prompts and responses are **logged securely** for traceability and quality assurance; access is limited to authorized personnel under least-privilege principles.

Specific AI functionality

Secure foundations for every AI capability

Simpplr integrates generative AI capabilities within secure, governed environments. Each capability leverages enterprise-grade providers and follows the same privacy, compliance, and security controls that apply across the platform.

CAPABILITY	PURPOSE	UNDERLYING SERVICE/ PROVIDER	DATA HANDLING & PRIVACY	SECURITY & CONTROLS
Knowledge Agent	Contextual Q&A and search using large language models.	Azure OpenAI or AWS Bedrock	• Search queries are processed and logged within Simpplr's environment for service analytics. • Logged data follows ISO 27001–compliant access-control and retention policies. No data is used to train public models.	• Encrypted in transit using SSL/TLS. • Logged data handled under ISO 27001–compliant access-control and retention policies.
EX Agent*	Conversational AI for employee and customer engagement.	Yellow.ai	• Context-restricted data processing only. • Yellow.ai does not use Simpplr customer data to train public models. • All interactions comply with GDPR, CPRA, and EU AI Act requirements.	• Encrypted in transit (TLS 1.2 +). • Simpplr monitors Yellow.ai for applicable security certifications and privacy standards.
Standard Search	Natural-language and semantic search within Simpplr.	Microsoft Azure OpenAI Service	• Uses Azure's zero-retention policy:ensuring data is never used to train or fine-tune models. All prompts and results are encrypted in transit using industry-standard SSL/TLS protocols. • Search queries are logged within Simpplr's environment for service analytics and search-quality improvements.	• Encrypted in transit (SSL/TLS). • All logged data follows ISO 27001-compliant access-control and retention standards.
Enterprise Search*	Federated and semantic queries across connected sources.	Microsoft Azure OpenAI Service	• Uses the identical Azure OpenAI Service pipeline as Standard Search, ensuring the same foundational data-protection standards. • Search queries are logged within Simpplr's environment for service analytics, consistent with Standard Search.	• Encrypted in transit (SSL/TLS). • All logged data handled under ISO 27001-compliant access-control and retention policies.
Core Intranet AI Agents (Writing Assistant, Summarization, Poll Generator, etc.)	In-platform assistive AI features that support content creation, summarization, and engagement.	Microsoft Azure OpenAI (deployed within Simpplr's infrastructure)	• Zero-retention architecture: prompts and results are not stored or reused for training. • Only contextually relevant data is processed.	• Encrypted in transit using SSL/TLS. • Operates within Simpplr's secure, ISO 27001–aligned infrastructure. All data adheres to compliant access-control and retention policies.

* denotes paid add-on functionality.

Configuration and governance

- Administrators can configure which AI functionalities are enabled or disabled within their environment.
- Simpplr works with third-party subprocessors for certain AI services and provides clear disclosure of their security, compliance, and certification obligations.
- Customers maintain ownership and control over all content and AI feature usage.

Compliance with regulations

Simpplr aligns its AI operations with global privacy, security, and AI governance requirements as outlined below:

FRAMEWORK / REGULATION	SIMPPLR ALIGNMENT
GDPR	• Privacy by design and default • Transparency via practices disclosed in privacy policy and Data Processing Agreement • Accuracy and Minimisation: data use limited to necessity • Security through appropriate technical and organizational safeguards
CPRA (California Privacy Rights Act)	Sensitive data protections; Automated decision-making rights retained by customers; Data minimization aligned to CPRA purpose limitation.
EU AI Act (2024/2025)	Risk-based approach: alignment with system classification; for high-risk systems: conformity assessments, documentation, and incident reporting planned; Transparency & human oversight applied.
EU NIS2 Directive (effective Oct 2024)	Aligning with obligations on 24-hour incident reporting, risk management, supply chain security, and resilience.
U.S. AI & Privacy Frameworks	Monitoring federal/state laws (including CPRA); following the White House AI Executive Order (2023).

AI governance & risk management

- **NIST AI RMF** alignment.
- **ISO/IEC 42001** adoption roadmap.
- **Incident response timelines:** GDPR **72 hours**; NIS2 **24 hours**.

Security certifications

Simpplr maintains recognized information-security and privacy certifications that underpin its AI and data-protection program:

CERTIFICATION / FRAMEWORK	SCOPE	STATUS
ISO 27001 (Certified against ISO/IEC 27001:2013)	Information Security Management System	Achieved
ISO 27701	Privacy Information Management	Achieved
SOC 2 Type II	Security, Availability, Confidentiality	Achieved (shared under NDA)
SOC 3	Public Report	Published
EU-U.S. & Swiss-U.S. DPF	Data Privacy Framework	Certified (via TRUSTe)
CSA STAR Level 1	Cloud Security Alliance	Achieved
CSA STAR Level 2	Advanced Attestation	Planned 2026
ISO/IEC 42001	AI Governance Standard	Planned 2026

Shared responsibility model

Simpplr's generative AI capabilities operate under a shared responsibility framework. This model defines clear accountability between Simpplr and its customers to maintain security, compliance, and trustworthy AI outcomes.

AREA	SIMPPLR RESPONSIBILITY	CUSTOMER RESPONSIBILITY
Infrastructure & platform security	Secure cloud hosting, patching, and encryption.	—
AI feature controls	Model configuration and subprocessor compliance.	—
Data privacy	Data isolation and retention/deletion processes.	—
AI feature usage	—	Enable only features compliant with company policy.
Employee use & AI output validation	—	Responsible for lawful use and interpretation of AI responses.

Security and compliance FAQ

Could the AI system generate confusion for some or all end-users or subjects on whether a decision, content, advice, or outcome is the result of an algorithmic decision?

Simpplr AI serves as an assistive, generative feature designed to enhance user experience, not to make autonomous decisions. All AI-driven responses are clearly presented as system-generated, and users retain full control over whether these responses are even shown to end users.

Did you establish any detection and response mechanisms for undesirable or adverse effects for the end-user or subject of the AI system?

Simpplr periodically reviews AI outputs against a sample dataset to assess accuracy, identify hallucinations, and mitigate undesirable behaviors. Any deviation or anomaly triggers internal evaluation and corrective updates.

Could the AI system have adversarial, critical, or damaging effects (e.g., to human or societal safety) in case of risks or threats such as design or technical faults, defects, outages, attacks, misuse, inappropriate or malicious use)?

The system operates within Simpplr's secure cloud infrastructure with robust access control, encryption, and monitoring. Azure OpenAI is used only for stateless inference, and no data is retained outside Simpplr's AWS-managed environment. Misuse risks are mitigated through role-based access, audit trails, and continuous security monitoring.

Did you put in place steps to monitor and document the AI system's accuracy?

Simpplr conducts periodic validation of AI responses using test datasets to ensure factual accuracy and minimize hallucination. These reviews are documented as part of the ongoing model evaluation and product quality process.

Did you consider the impact of the AI system on the right to privacy, the right to physical, mental, and/or moral integrity, and the right to data protection?

All AI-related processing complies with global data protection standards (including GDPR). Data used for inference is anonymized and transient. No personal data is shared or stored within Azure OpenAI, and all content remains within Simpplr's AWS environment under strict access controls.

Is your AI system being trained, or was it developed, by using or processing personal data (including special categories of personal data)?

The system does not use customer or user data for training. All Azure OpenAI interactions are inference-only and ephemeral; no data is retained or reused.

Did you put adequate logging practices in place to record the AI system's decisions or recommendations?

Prompts and responses are logged for traceability and quality monitoring. Logs are securely stored and accessible only to authorized personnel for support, audit, and improvement purposes.

Do users know when they are interacting with AI rather than a human (e.g., in chatbots)?

Simpplr clearly identifies AI-generated responses within the user interface.

Did you establish a strategy or a set of procedures to avoid creating or reinforcing unfair bias in the AI system, both regarding the use of input data as well as for the algorithm design?

Because embeddings and inferences are generated from customer-provided content, bias is minimized through content neutrality. Simpplr performs periodic reviews to identify and correct potential bias or factual inconsistencies.

Did you assess whether some groups could be disproportionately affected by AI outputs?

The system is designed to provide equitable access and uniform results across all users. There is no differential treatment based on personal attributes, roles, or demographics.

Does the AI system impact human work or job arrangements?

The AI system augments existing workflows by improving information retrieval and productivity.

Did you ensure that workers understand how the AI system operates, which capabilities it has and which it does not have, ensuring proper understanding of its assistive role?

Product documentation, in-app guidance, and customer training sessions explain the system's purpose, boundaries, and operating model.

Did you establish mechanisms that facilitate the AI system's auditability (e.g., traceability of the development process, the sourcing of training data, and the logging of the AI system's processes, outcomes, positive and negative impact)?

The system's data flow, processing steps, and API interactions are fully documented. Logging and internal review processes provide traceability for audit and compliance purposes.cpport third-party or customer audits where required under regulatory obligations.

About Simpplr

Simpplr is the AI-powered platform that unifies the digital workplace – bringing together engagement, enablement, and services to transform the employee experience. It streamlines communication, simplifies interactions, automates workflows, and elevates the everyday experience of work. The platform is intuitive, highly extensible, and built to integrate seamlessly with your existing technology. More than 1,000 leading organizations – including AAA, the NHS, Penske, and Moderna – trust Simpplr to foster a more aligned and productive workforce. Headquartered in Silicon Valley with global offices, Simpplr is backed by Norwest Ventures, Sapphire Ventures, Salesforce Ventures, and Tola Capital. Learn more at simpplr.com.