cicero

Cicero: what your IT, AI, and security teams need to know

We know you're focused on finding innovative solutions that deliver value while maintaining the highest standards of security and compliance. Cicero is a secure and compliant AI-powered platform for realistic roleplay training, designed to meet your organization's needs without compromising data privacy or security.

This guide is tailored to help you share critical information about Cicero with your IT, AI, and security teams so they can evaluate our solution with confidence. It addresses key security and data privacy considerations, demonstrating how Cicero protects sensitive information and aligns with industry best practices. Below, we've addressed the most common questions and concerns your technical stakeholders may have.

How Cicero **protects your data**

Data minimization

- Only the minimum data required for roleplay scenarios is gathered.
- Fully compliant with GDPR, PIPEDA, and other global privacy regulations.

Data encryption

- All data is encrypted both at rest and in transit using industry-standard protocols to prevent unauthorized access.

Vector database

- Data is stored securely in a Vector Database, transformed into binary representations for use in scenarios.
- This data is not used to train large language models (LLMs), ensuring confidentiality.

Data lifecycle management

- Default retention: contract duration + 30 days.
- US/Canada standards: up to 7 years (configurable by your organization).
- Trial data: deleted within 30 days of written notification if you choose not to proceed.

Access control

- Role-based access control (RBAC) and multi-factor authentication ensure only authorized personnel can access sensitive data.

Data residency

- Data is stored in the US East-1 region by default, unless otherwise specified, offering regional data residency options to comply with specific regulatory requirements.

Key questions **your IT, AI, and security teams may have**

What data does Cicero collect?

- Only minimal data necessary for roleplay scenarios. No unnecessary personal information is tracked.
- Benefit: reduces risk and emphasizes our privacy commitment.
- Includes essential information like name, email, and optional learning progress tracking.

How is data handled and what is used to train Cicero?

Cicero is designed to protect your data and ensure privacy at every step.

- Uses a Retrieval-Augmented Generation (RAG) backend for scenario creation.
- Your documents are never used to train, fine-tune, or embed large language models (LLMs).
- No client data or prompts are shared or used to train models for other clients.
- No external consumer data is used; all outputs are based solely on your supplied documents and evaluation criteria.

Does end-user input train the AI?

- No. Your input is not used to train or improve the LLM. It is solely used for evaluating user performance within specific roleplay scenarios.

What happens to trial data?

- All trial-related data will be securely deleted within 30 calendar days of written notification if you decide not to move forward.

How does Cicero adapt to emerging threats?

- Continuous monitoring by our development team and security experts ensures proactive vulnerability management through regular updates, penetration testing, security audits, and real-time threat intelligence.

What measures prevent exploitation or attacks?

- Input sanitization and filtering mechanisms prevent malicious inputs.
- Data encryption protects information at rest and in transit.
- Adversarial testing ensures resilience against vulnerabilities.
- Continual human oversight enhances security monitoring.

How does Cicero integrate with our systems?

- Cicero integrates seamlessly with systems like LMS or HR platforms via APIs or Single Sign-On (SSO). Detailed integration documentation is available upon request.

What are the legal and contractual considerations?

- Standard Data Processing Agreements (DPAs) are available to ensure compliance with relevant privacy regulations. Our legal team can address specific contractual requirements.

Architecture, integration, and operational resilience

Can you provide a high-level architecture and network security overview?

- Yes, Cicero can provide a detailed architecture diagram under NDA. The platform is built on AWS with enterprise-grade security, network segmentation, and full encryption of data flows.

What integrations does the application support?

- SSO and SCORM-compliant LMS integration supported; connections can be bidirectional if required. APIs are available.

What data transmission methods are used?

- Fully encrypted HTTPS/TLS is standard. The SaaS model is consumption-based and scales based on user numbers.

What is your hosting and availability model?

- Highly available AWS-hosted servers with redundancy.

Can you describe your disaster recovery (DR) scenario?

- Optional second cloud provider as part of continuity/disaster recovery planning. Backups and restore tests performed regularly.

How does the platform support scale-up and scale-out scenarios?

- Scales dynamically with usage. Uploaded documents are securely extracted and stored in a PGVector database, with 30-day deletion for GDPR compliance.

How is data stored and encrypted?

- Data and documents are encrypted at rest using AES-256 and stored securely in AWS S3, accessible only via the application.

What cryptographic frameworks are used?

- Uses industry-standard AES-256 encryption and secure key management.

What type of events are logged and monitored?

- Authentication events, system errors, data access requests, administrative actions, and security events are logged and monitored in real-time.

Additional **questions and answers**

Is Cicero privacy compliant?

Cicero complies with GDPR and PIPEDA and continually improves data protection.

How is my data protected

We utilize multiple security layers:

- Strong encryption: data is scrambled in storage and transit (think digital lockbox).
- Cloud protection: secure cloud with built-in defenses against cyberattacks and malware (constantly vigilant security team).
- Data organization: data is separated, with dedicated resources available (organized compartments).

Secure locations: 0ur offices meet high security standards, like high-security buildings (ISO 27001 & SOC 2).

Is client data or external data used to train Cicero's AI model?

No. Cicero never uses client data or external consumer data to train its AI models.

- Content generation is based only on your supplied documents and criteria.
- Sensitive information remains private and is not added to any shared training set.

What is Cicero's privacy by design approach?

Privacy is integrated from the start:

- Data interoperability: secure data sharing across platforms.
- Enhanced trust: robust privacy builds confidence.
- Regulatory compliance: adherence to data privacy regulations.
- Improved collaboration: secure data sharing across departments and partners.
- Cost efficiency: reduced infrastructure costs.

What policies and procedures govern Cicero's AI lifecycle and how is bias managed?

Cicero follows strict governance and oversight to ensure fairness and compliance.

- Human-in-the-Loop (HITL): All roleplay scenarios require 100% human sign-off before deployment.
- Comprehensive testing includes controlled user scenarios, edge case stress tests, and benchmarking.
- Bias mitigation techniques are applied; no user data is used for model training.
- Transparent reporting: Admin dashboards, self-reporting tools, and enterprise-level score reports.
- Data retention policies govern document and metrics storage throughout engagement.

How does Cicero integrate and evolve with learning systems?

Easy integration: works seamlessly with existing systems like Articulate Storyline 360™.

- Leveraging SCORM: ensures compatibility with SCORM-compliant systems.
- Interactive content creation: enables engaging learning experiences through its Storyline plugin.
- Detailed progress tracking: utilizes xAPI to monitor learner performance effectively.
- Secure and efficient integration: guarantees data safety and smooth implementation across any compliant LMS.

Continuous improvement: constant security enhancements and pursuit of higher certifications (ISO 27001 in 2025).

How does Cicero ensure transparency, explainability, and communicate product updates?

Cicero keeps you informed and in control.

- Product updates are shared via Customer Success briefings, monthly webinars, portal banners, and in-app notifications.
- Explainability: Whitepapers, research, and (under NDA) block diagrams are available.
- Three core components: Creator, Scenario, Evaluator-each with clear, human-readable results and actionable feedback.
- Scoring and evaluation are transparent and configurable by administrators.

Key security features at a glance

FEATURE	DESCRIPTION
Encryption	Data encrypted at rest and during transmission using industry standards.
Vector database	Securely stores binary-transformed data; does not train LLMs.
Customizable retention	Cicero clients define retention periods for compliance needs.
Adversarial testing	Ensures resilience against malicious inputs or vulnerabilities.
Real-time threat monitoring	Logs potential threats for rapid mitigation using advanced threat intelligence tools.
Input sanitization	Filters and sanitizes all inputs before processing by the system.
Access control	Role-based access control with multi-factor authentication (MFA).
Incident response	Defined processes for handling security incidents or breaches effectively.

Why your IT, AI, and security teams can trust Cicero

01

No training on your data:

your input remains private—it's never used to train our language models.

02

Customizable retention policies: . you control how long data is stored based on your organization's requirements.

03

Proactive threat management: continuous monitoring ensures rapid adaptation to emerging threats.

04

Transparency: **clear communication** about how your data is handled every step of the way.

05

Bias mitigation: techniques are employed to identify and reduce potential biases in AI models.

06

Model explainability: decisions are grounded in understandable logic that aligns with user expectations.

Client testimonials Case studies:

Medtronic

Medtronic

The Challenge: Medtronic, a global leader in medical technology, sought to transform its training programs for its virtual surgery platform in North America. The goal was to improve performance and adoption of the platform while adhering to the highest standards of security and compliance within the healthcare industry.

The solution: Cicero's AI-Powered Training Platform

Medtronic used Cicero to implement a cutting-edge AI and XR-powered training program. Cicero's platform offered a flexible deployment model, utilizing a hybrid hosting solution with the frontend on CGS's Immersive public cloud and the backend on Medtronic's private cloud. This allowed for scalability and accessibility while maintaining strict control over sensitive data. Alternatively, Cicero could have been deployed on Medtronic's public or private cloud.

The results:

- Enhanced training performance: by leveraging AI and XR, Cicero created immersive and engaging training experiences that accelerated the learning process and improved user proficiency with Medtronic's virtual surgery platform.
- Robust security and compliance: Cicero's platform was designed with security as a top priority, meeting the stringent compliance requirements of the healthcare industry.
- Industry recognition: Medtronic's innovative training program, powered by Cicero, was recognized as First Runner-Up in the IDC Future Enterprise Awards for Best in Future of Work.

> "Cicero has transformed our training programs while maintaining the highest standards of security and compliance."
>
> –Director of Training, Medtronic North America

cicero

Next steps

Share this guide with your IT, AI, and Security teams as a starting point for evaluation. For a technical deep-dive, demo, or further information, contact cicero@cgsinc.com

Let’s work together to bring secure, innovative roleplay solutions to your organization!