Civinc Security Policy

Version 2.0 · Effective 29/7/2026 · Reviewed annually

This document describes the technical and organisational measures we take under Article 32 GDPR. It forms part of our Data Processing Agreement. Where this document and that agreement differ, the agreement prevails.

1. What we protect, and how little there is

The strongest security measure is not collecting the data in the first place. Civinc never asks participants for a name, email address, telephone number or other information linking to their personal identity. Session data is attached to a randomly generated identifier, and we hold no key linking that identifier to a person. A breach of our systems would not reveal who said what, because we do not know.

Chat content is the most sensitive data we hold, and we hold it for the shortest time: it is deleted permanently within 72 hours of the end of a session. No Civinc employee has access to it at any point.

2. People and access

- Back-end access to systems holding user data is limited to director-level personnel. No one else has it, and it is granted only where the role requires it.
- Everyone with such access is bound by a written confidentiality obligation that survives the end of their engagement.
- Access rights are revoked and credentials invalidated as part of our offboarding process when someone leaves.
- Access levels across all systems are reviewed at least once per year.
- Administrative accounts on our own systems and on those of our providers are protected with multi-factor authentication.
- Secrets and API keys are held in a managed secret store, never in source code, and are rotated when personnel change.

3. Infrastructure

Civinc runs entirely on Google Cloud Platform, on servers located in the European Union. We operate no servers of our own.

- All data is encrypted in transit over HTTPS/TLS, and encrypted at rest.
- Customer data is logically isolated.
- We use Firebase App Check, which verifies that calls to our back end come from our own application and not from a third party.
- Google Cloud holds ISO 27001, ISO 27017, ISO 27018 and SOC 1, 2 and 3 certification, and runs continuous penetration testing of its infrastructure.

Those certificates cover the infrastructure we build on, not our own organisation. The next section sets out where Civinc itself stands.

4. Certification

Civinc held ISO 27001 certification until June 2026. We chose not to renew it for feasibility reasons related to being a small enterprise. What we kept is the management system underneath the certificate. Our access control, supplier management, incident response, risk assessment and annual review processes are the ones we built to meet ISO 27001, and they remain in operation. They are described throughout this document.

5. Data handling and retention

Data	Retention	Protection
Chat content	Maximum 72 hours after the session	Encrypted at rest, no employee access, permanently deleted. Held only so a participant can recover their conversation after a page refresh and so any AI analysis can run.
Answers and session metadata	Term of the customer agreement, maximum three years	Pseudonymous. Deleted earlier on the customer's instruction.
Administrator accounts	Term of the customer agreement	Name, email address and organisation only.
IP addresses	Not stored by Civinc	Processed by Firebase to establish the connection; Civinc has no access to them.

Deletion is permanent and applies to backups within the same cycle. On termination of a customer agreement we delete or return all data belonging to that customer within thirty days, at the customer's choice.

6. AI processing

Where a customer enables AI analysis, conversation fragments are sent to Mistral AI SAS, a French company, and processed on servers in the European Union. Mistral acts as our processor under a data processing addendum.

- Fragments are detached from the session context and isolated from all other data. Mistral does not receive the session as a whole, the customer's identity or the user identifier.
- Data sent to Mistral is not used to train its models. This opt-out is set at organisation level.
- Experimental ("Labs") models are disabled, because they permit training regardless of that opt-out.
- We do not use the feedback endpoints, which would place submitted content outside the opt-out.
- Mistral retains fragments for a maximum of 30 days for abuse monitoring, after which they are permanently deleted.

AI analysis is optional per session, and the analysis only runs once enough messages have been sent across enough separate conversations that no individual contribution is distinguishable. Output never contains direct quotes.

7. Secure development

Security is a standing item in our product process, not a phase at the end.

- Code conventions enforced automatically through static analysis.
- Established frameworks used to protect against common attack classes, including XSS, CSRF and injection.
- Continuous integration with unit and integration tests; every release is tested before it reaches production.
- Dependencies monitored continuously and kept current; security advisories are acted on as they arrive.
- Continuous vulnerability scanning with dedicated third-party tooling.
- Changes are peer-reviewed before they are merged.

8. Monitoring, backups and continuity

- Access to production systems and administrative actions are logged.
- Backups are made based on pre-established frequencies.
- Because chat content is deleted within 72 hours, it is by design not present in older backups.

9. Incidents and breaches

Security incidents can be reported to hello@civinc.co at any time, by anyone. Reports reach our CTO directly.

- Every incident is documented, assessed and assigned an action plan including mitigating measures.
- Where a personal data breach affects a customer, we notify that customer without undue delay and in any event within 48 hours of becoming aware, with the detail required under Article 33(3) GDPR.
- We do not notify supervisory authorities or data subjects on a customer's behalf unless instructed in writing, since that decision belongs to the controller.
- Where individuals are directly affected and reachable, we contact them through the appropriate channel.
- After resolution we review what happened and what should change.

10. Sub-processors

Provider	Entity and location	Role
Google Cloud Platform	Google Cloud EMEA Ltd — EU region	Hosting, storage, database and analytics infrastructure.
Google Firebase	Google Cloud EMEA Ltd — EU region	Connection between the participant's device and the platform.
Mistral AI	Mistral AI SAS — Paris, France	AI analysis of conversation fragments, where enabled by the customer.

Each sub-processor is bound by a written agreement imposing data protection obligations no less protective than our own. We inform customers before adding or replacing a sub-processor, so they can object.

Data is stored within the EEA. Our hosting provider belongs to a group with entities outside the EEA whose personnel may access data remotely for support and maintenance; under EDPB guidance that access is a transfer, and it is covered by the EU-U.S. Data Privacy Framework with Standard Contractual Clauses as a fallback. Our AI provider processes exclusively within the European Union.

11. Review

Activity	Frequency
Revoke system, hardware and document access	On departure
Review access levels across all systems and personnel	Annually
Review this policy and the sub-processor list	Annually, and on any material change
Keep critical libraries and dependencies current	Continuous
Unit and integration testing	Continuous
External vulnerability scanning	Continuous

Questions about this policy: hello@civinc.co