CyberResilient

SECURITY OVERVIEW

# Security at CyberResilient

How the platform is built, where your data lives, and how we work with information security as a company.

CyberResilient AB (556702-0200) · Version 1.2 · 15 September 2026 · Classification: Public

In CyberResilient you collect assessment answers, evidence, an action plan and a risk register. Together they form a detailed picture of where your organisation is vulnerable. This document describes how we protect that picture.

We only describe what is in place today. More about us, and documents you can request, is at cyberresilient.se/en/trust.

## 1 What data does the platform handle?

**You enter three kinds of data.** Details about your organisation when you get started, such as name, organisation number, sector and size. Answers to the assessment questions, with any notes you write alongside them. And evidence: documents you upload to support an answer.

**The platform scores and writes.** Answers are scored against the framework you chose, and a report is produced with observations and recommendations. The recommendations become objectives and controls in the action plan, and the risk register points back to the same controls.

**Data leaves the platform mainly through three flows.** Text from evidence, answers and notes, and a limited part of the organisation profile, goes to our AI provider (section 4). Email from the platform is sent through our email provider. Billing details go to our payment provider. We do not sell your data.

## 2 Where does your data live?

The platform is hosted by Hetzner in Falkenstein, Germany. The production server runs the application, the database and background jobs. Uploaded evidence is stored in Hetzner Object Storage in the same region. The platform's hosting, storage and AI processing take place within the EU. Email and payment suppliers may transfer data outside the EU, see the privacy policy.

| WHAT | WHERE |
|---|---|
| **Platform and database** | Hetzner, Falkenstein, Germany |
| **Uploaded evidence** | Hetzner Object Storage, same region |
| **AI processing** | Berget AI, Sweden |

**Exposure.** The production server has only ports 80 and 443 open to the internet. The database and cache publish no ports and cannot be reached from outside. All traffic to the platform uses HTTPS.

**Our own access.** No server has public SSH. Operational access runs over an encrypted network, with personal accounts and two-factor authentication, governed by the network's access rules. All our staff use two-factor sign-in on every internal system.

**Build and backups.** Servers, networks and firewalls are described in code with Terraform, which makes the environment reviewable and reproducible. All servers are backed up automatically every day. Production and test run on their own servers with their own databases.

A full list of subprocessors is available on request.

## 3 How is access protected?

### Sign-in

- Passwords must be at least twelve characters and pass a strength check that rejects common and guessable passwords. They are stored as Argon2id hashes (64 MiB memory, three iterations, four threads).
- Two-factor authentication with one-time codes (TOTP, RFC 6238) and ten recovery codes. The secret is stored encrypted; recovery codes are hashed and encrypted.
- Sign-in, sign-up, one-time codes, password reset and the two-factor step are rate limited. After ten failed attempts the account is locked and all sessions are ended.
- Sign-in and sign-up respond the same way whether or not an address exists, so nobody can find out which accounts exist.
- Each session has a random 256-bit token that is stored only as a hash. Sessions end after inactivity and after a maximum lifetime, and can be ended manually.

### Separation between organisations

The organisation a request applies to is resolved against the database every time, and requires both an active membership and an active role. The cookie that carries the choice is signed and treated only as a hint.

When a partner works in your organisation, three things are required at once: an active membership at the partner, an active role there, and that your organisation is linked to that specific partner. The partner's access can never exceed its role. Deleting the organisation requires direct ownership and a fresh two-factor check, and is blocked entirely for partners.

### Encryption and evidence

On top of encryption in transit, sensitive fields are encrypted in the application with AES-256-GCM and versioned keys, so older data stays readable after a key rotation. This covers two-factor secrets and recovery codes, IP addresses and browser details in sessions, security logs, email tokens and the processed results a report is built from.

Uploaded evidence is never served directly. Every access passes a checkpoint that checks permission, logs the access and issues a signed link valid for five minutes. Files are checked on upload: file type, size, and that the content is what the file claims to be.

### Traceability

Every change to the risk register is stored as its own event with who, what and when, in a log the application never changes or deletes from. Access to evidence is logged, as are edits to a completed assessment, with the value before and after.

## 4 How do we use AI?

The platform uses language models to read uploaded evidence, link it to the right question, and write the report's observations and recommendations. Recommendations are advisory. The platform makes no automated decisions about individuals.

**One provider, in Sweden.** We use a single AI provider, Berget AI. Every model we run has open weights: *Mistral-Small-3.2-24B-Instruct-2506* for analysis and reports, with *gemma-4-31B-it* as fallback,

*multilingual-e5-large* for embeddings and *bge-reranker-v2-m3* for ranking results. Gemma is developed by Google, but like all our models it runs at Berget. No call goes to Google or any other US model provider.

**What is sent.** Text from your evidence, your answers and your notes on the questions. From the organisation profile, three fields are sent: sector, size band and the business description you wrote yourselves. Organisation number and turnover are never sent.

**No training.** We do not train or fine-tune any model on customer data. Prompts and responses containing customer data, and embeddings, are classified as confidential, at the same level as customer data and source code.

## 5 How do we build the platform?

**Review.** Our secure development guideline requires another developer to review every change before it is merged, and sets the OWASP Top 10 as the minimum standard for web and API code.

**Automated tests.** Every change runs through linting, type checking, format checks and the full test suite against a real database and cache. Test coverage is measured against a set threshold.

**Deployment.** Deployments run automatically from version control, never by hand on a server. A deployment only counts as successful once the application actually responds.

## 6 How do we govern security as a company?

Since 20 July 2026 we have had a documented information security management system, approved by our CEO. It consists of a policy, three guidelines, two instructions, two procedures, and templates for risk analysis, statement of applicability and follow-up. The structure follows ISO/IEC 27001 Annex A, and our risk work follows a method based on ISO/IEC 27005.

**Certification is the goal.** We have not set a date and do not claim to be certified.

**Classification and confidentiality.** Customer data, source code, secrets and logs are confidential by default. Employees and contractors are bound by confidentiality agreements.

**What the management system commits us to:**

| AREA | COMMITMENT |
|---|---|
| **Incidents** | First response within 1 hour for the most severe, within 4 and 24 hours for the next levels |
| **Customer notification** | As agreed, and where there is no agreement, within 24 hours of confirmed impact |
| **Personal data breaches** | Assessment of notification to the Swedish Authority for Privacy Protection (IMY) within 72 hours |
| **Suppliers** | Classified by criticality. AI providers are always critical and are reviewed at least annually |
| **Follow-up** | Management review and internal audit of the management system every year |

**Incident management.** Roles are assigned for incident lead, technical handling and customer communication, with decision support for each incident type, including a specific procedure for data leaking through an AI service.

**Questions?** Governing documents, security questionnaires or anything else: security@cyberresilient.se · cyberresilient.se/en/trust
**Found a vulnerability?** Send what you found and how to reproduce it to security@cyberresilient.se, also listed in our security.txt. Do not test against other customers' data.