Security Statement

Kanriva™ operations platform and FlowCrew™ worker app

Version 1.0 · 29 June 2026 · Kaizen Kanri (KVK 42049431)

1. Purpose of this document

This statement describes how Kaizen Kanri hosts, protects and deletes the data you put into Kanriva™. It is written to be handed to an IT manager, a privacy officer or a works council without translation, and to replace the usual vendor security questionnaire — section 12 answers those questions directly.

Everything in this document is a statement of what is true today. Where a control is not yet in place, it is listed in section 11 rather than omitted. If you find a claim here that you cannot verify during a walkthrough, treat it as a defect and report it to info@kaizenkanri.nl.

2. What the platform processes

Kanriva™ plans, publishes and measures work. To do that it holds workforce data: names and employee numbers, department and function, contract hours, shifts and availability, absence and holiday, presence and actual times, productivity per person and per process, documented follow-up of formal conversations, uploaded documents such as certificates or damage photographs, and shift-linked messages and handovers.

That is personnel data, and it is treated as such. Two consequences:

- **Special categories are neither required nor wanted.** Absence can be recorded as a type or code. Medical details, diagnoses and occupational-health reports must not be entered; your users are contractually instructed not to.
- **Performance data is access-controlled.** Productivity and personnel screens are separate modules with their own roles, so a shift planner does not automatically see personnel records.

3. Hosting and data location

The application, the database and uploaded files run on infrastructure provided by **Strato AG in Germany**, inside the European Union, on a managed container platform operated together with our deployment partner. No workforce personal data is stored outside the European Union.

Three narrow exceptions are named rather than buried:

- the **payment provider** processes billing contact and payment data of your administrators, and never workforce data;
- the **transactional email provider** processes the recipient address and the content of platform emails;
- **mobile push delivery**: where push notifications are enabled on an Android device, hand-off to the device is performed by Google's notification service, and notification text is kept free of personnel details. Notifications inside the platform are generated and stored on our own EU infrastructure.

The first two are engaged under the European Commission's Standard Contractual Clauses. The complete list, with roles and locations, is in the **Sub-Processor Register**.

Everything else stays in the EU, including application logs, error logs, the audit trail, uploaded files and backups.

4. Encryption

- **In transit** — all browser and mobile traffic is encrypted with TLS. Plain HTTP is not served. Authentication cookies are marked Secure.

- **At rest** — the database, uploaded files and backups sit on provider-managed encrypted storage. To be precise, because a security reviewer will ask: this is **storage-level encryption with provider-managed keys**, not per-field application encryption, and not customer-managed keys.

5. Accounts, sessions and access

- Passwords are hashed with **bcrypt** and individually salted. They are never stored or logged in plain text, and cannot be read by Kaizen Kanri.
- Sessions use signed JSON Web Tokens delivered in **HttpOnly, Secure, SameSite** cookies, valid for a maximum of seven days, after which the user re-authenticates.
- Invitation and password-reset links are single-use and time-limited.
- Access is **role-based and enforced on the server** for every request. Hiding a button is never the only control.
- You administer your own users and roles. No customer role can reach another customer's data.
- One exception, stated openly: a **named Kaizen Kanri master-admin account** exists for tenant administration and support. Its use is limited to support, incident resolution and maintenance, and it is bound by the confidentiality and least-privilege terms of the DPA. Infrastructure-level access (hosting console and database tooling) inevitably sits with the operator, as it does with any hosted platform; it is not a separate back door, and it is covered by the same obligations.
- **Multi-factor authentication is not yet available.** See section 11.

6. Separation between customers

Every customer is a separate tenant. Every database query, export, report and background job is scoped to the authenticated user's tenant on the server side. There is no shared list, no global search and no cross-customer reporting.

This is the control we take most seriously, because it is the one that would matter most if it failed. It is therefore covered by **automated regression tests** that call the platform as one tenant and assert that no record belonging to another tenant is ever returned. The covered surfaces are the REST API and its list and detail endpoints, Excel and PDF exports, file downloads, search, reports and dashboards, background jobs, and the FlowCrew mobile endpoints. Those tests run as part of the release process, and the test report can be shared on request.

The claim is deliberately scoped: isolation is enforced in the application layer on every request, and verified by tests on the surfaces listed above. It is not a claim about the hosting provider's own infrastructure tooling, which is covered by the administrative route described in section 5.

7. Traceability

The platform keeps a change history of operational records: who changed what, when, and what the previous value was. Schedules, workload figures, presence corrections, handovers and personnel entries are therefore reconstructable, which is what makes the data usable in a formal conversation or a dispute.

Application errors are logged for troubleshooting only. Logs are not a monitoring tool and are not used to observe individual employees.

8. Backups and recovery

Full backups are taken **monthly** and stored within the European Union. Restores are tested when a change affects the data model.

Be aware of what that means: the recovery point objective is **up to one month**, and there is no point-in-time recovery. For most planning operations this is acceptable because the schedule is rebuilt weekly and exports are held locally, but if your risk assessment requires daily backups, say so before signature — it can be arranged per

customer. Details are in the **Backup, Retention and Exit Statement**.

9. Deletion, export and exit

- You can export from every module to Excel or PDF at any time, without asking us. That is deliberate: your data should never be hostage to a subscription.
- Deleting an account archives it and keeps it recoverable for **thirty days**, so an accidental or disputed cancellation does not destroy your records. After thirty days an automated process permanently removes the tenant and its records from the production database.
- On written request we provide a complete structured export before deletion, and confirm in writing when the last backup containing your data has been overwritten.

10. Artificial intelligence — what we do not do

The staffing recommendation engine is **deterministic and rule-based**. It calculates required hours from the workload you enter, compares them with availability, contract hours and shift rules, and proposes a staffing per department. Every proposal is explainable, reproducible, and confirmed by a planner before it becomes a schedule.

Consequently:

- **No third-party AI or large-language-model service processes your data.** No employee record, message or productivity figure is sent to any AI provider.
- **No model is trained, fine-tuned or evaluated on your data**, by us or by any sub-processor.
- **No production data is used for benchmarking, analytics or marketing.** Demonstration environments, including the public demo tour on our website, contain invented names and invented figures only.

If a supplier cannot tell you plainly where employee data goes when they say "AI", that is the question worth asking. Here, it does not leave the platform.

11. What is not in place — stated openly

- **Multi-factor authentication** for platform accounts is not yet available. It is on the roadmap. Compensating controls: strong unique passwords, seven-day session expiry, prompt removal of leavers, role-limited access to personnel modules.
- **No external penetration test** has been commissioned, and we hold **no ISO 27001 certificate and no SOC 2 report**. We will not imply otherwise to win a tender.
- **Backups are monthly**, not daily, with no point-in-time recovery.
- **Kaizen Kanri is a small, owner-operated company.** Continuity is addressed by version control, written architecture and handover documentation per module, managed EU hosting that runs independently of any individual, and data that you can export at any time. The continuity statement on the website describes this in full.
- **No formal ISMS or certified audit programme.** What exists instead: this document, the sub-processor register, the incident response procedure, automated isolation tests, and a standing offer to let your IT lead inspect the setup during a walkthrough.

12. Vendor security questionnaire — pre-answered

Company and contract

#	Question	Answer
1	Legal entity and registration	Kaizen Kanri, Netherlands, KVK 42049431

#	Question	Answer
2	Role under GDPR	Processor. The customer is the controller.
3	Is a DPA available?	Yes — downloadable immediately in English and Dutch, counter-signed on request before any employee data is entered. Your own template is also acceptable for review.
4	Data protection contact	info@kaizenkanri.nl
5	Is a DPO appointed?	No statutory obligation to appoint one; the owner is the accountable contact.
6	Sub-processors disclosed?	Yes — Sub-Processor Register, with 30 days' notice of changes and a right to object.

Data and location

#	Question	Answer
7	Where is data hosted?	Strato AG, Germany (EU).
8	Is data ever outside the EU?	Workforce data: no. Three exceptions: billing contact data (payment provider) and platform email content (email provider), both under Standard Contractual Clauses, and Android push hand-off through Google's notification service with no personnel details in the notification text.
9	Categories of personal data	Identity, employment, scheduling, absence, presence and actuals, productivity, documented follow-up, uploaded documents, app account data, shift-linked messages, audit metadata.
10	Special category data	Not required and not permitted; users are instructed not to enter it.
11	Data retention	Term of the subscription plus a 30-day recovery window, then permanent deletion.
12	Data portability	Excel and PDF exports in every module, self-service; full structured export on request.

Security controls

#	Question	Answer
13	Encryption in transit	Yes, TLS; HTTP not served; Secure cookies.
14	Encryption at rest	Yes — storage-level encryption with provider-managed keys, covering database, uploaded files and backups. Not per-field application encryption; no customer-managed keys.
15	Password storage	bcrypt with per-password salt.
16	Session management	Signed JWT in HttpOnly/Secure/SameSite cookie, maximum seven days.
17	Multi-factor authentication	Not yet available; on the roadmap.
18	Single sign-on / SAML	Not available today.

#	Question	Answer
19	Role-based access control	Yes, granular per module, enforced server-side.
20	Customer separation	Separate tenant per customer; every application query scoped server-side; automated regression tests cover API, exports, file downloads, search, reports, background jobs and mobile endpoints.
21	Audit trail	Yes — who changed what, when, and the previous value.
22	Admin access to customer data	A named master-admin account for tenant administration and support, least privilege, bound by the DPA. Infrastructure-level access sits with the operator, as with any hosted platform.
23	Penetration test	None commissioned to date.
24	ISO 27001 / SOC 2	Not held.
25	Vulnerability management	Dependencies updated on release; security-relevant updates applied promptly.
26	Anti-virus / endpoint	Standard endpoint protection on development machines; server workloads are managed containers.

Operations

#	Question	Answer
27	Backup frequency	Monthly full backups, stored in the EU.
28	Recovery point / time objective	RPO up to one month; RTO target one business day. A shorter RPO can be agreed per customer.
29	Restore testing	Tested when a change affects the data model.
30	Disaster recovery plan	Redeploy from version control onto managed hosting and restore the latest backup; documented per module.
31	Uptime commitment	Continuous availability while hosting infrastructure is operational. No SLA percentage on trials/pilots; SLA available on annual agreements.
32	Planned maintenance	Announced in advance; scheduled outside shift-change peaks where possible.
33	Breach notification	Within 48 hours of becoming aware, in writing, per the Incident Response Procedure.
34	Incident response plan	Yes — published document, reviewed annually.
35	Support channel and response time	info@kaizenkanri.nl; reply within one business day, usually faster during CET business hours.

Development and AI

#	Question	Answer
36	Source control and change history	Yes, every change attributable.

#	Question	Answer
37	Automated testing	Yes — authentication, tenant isolation, planning logic and exports.
38	Production data in test environments	No. Test and demo environments use invented data.
39	AI / LLM processing of customer data	None. The staffing engine is deterministic and rule-based, running on our own infrastructure.
40	Customer data used for model training	Never, by us or by any sub-processor.
41	Third-party analytics or tracking in the app	None in the application.
42	Mobile app distribution	FlowCrew™ on Google Play (com.flow_crew). iOS under consideration.

If your questionnaire contains a question that is not answered above, send it to info@kaizenkanri.nl. It will be answered in writing, including the awkward ones.

Kanriva™ and FlowCrew™ are trademarks of Kaizen Kanri.