

Security & LLM Data Use — Compliance Statement

A candid, audit-facing description of what Quorval does, what it does not do, and where our controls sit relative to GDPR (EU/UK), APPI (Japan), CCPA/CPRA, HIPAA, PIPEDA and LGPD.

Owner: CreativePrinTech LLC (dba Quorval) | Version 1.1 | Effective August 02, 2026 | Classification: Public

Purpose. This document helps auditors, DPOs, procurement and security reviewers make an informed decision. Each control is described with its current status — implemented, partial, planned or delegated to a sub-processor — and certification status is stated factually throughout.

Change from version 1.0 (26 July 2026): this revision reports control changes made since publication — multi-factor authentication is now enforced for privileged roles, every private storage bucket is ownership-scoped, tenant isolation is verified by deterministic automated probes rather than by assertion, all recurring automations are instrumented and watchdogged, a control register with scheduled evidence is operating, and our own AI systems are evaluated on a schedule. Section 6 and the new sections 7 and 8 carry the detail, including a suite currently below its target threshold.

1. What Quorval is, in plain terms

Quorval is a SaaS workspace for business analysts, project managers and their stakeholders. Users create projects, plan work using BABOK-aligned phases, capture meetings (including with an optional recording bot), generate artifacts — requirements, risks, process maps, change communications — and collaborate with internal teammates and external guests.

The data we process on behalf of a customer includes user account information (name, email, role, organization), project metadata, meeting transcripts and recordings when the user opts in, uploaded documents, requirements, risks, decisions, tasks, chat messages with the AI coach, and audit and e-signature evidence.

2. Roles under GDPR / APPI and equivalent laws

For content a customer organization uploads or generates in its workspace, Quorval acts as a data processor (GDPR Art. 28) and as a trustee of entrusted personal information (APPI Art. 25). The customer organization is the controller and is responsible for lawful basis, notice and data-subject requests toward its own end users and stakeholders. For the account of the signed-up user themselves — billing email, login, telemetry, marketing — Quorval is the controller.

Sub-processors. The following list is authoritative and will be updated with 30 days' notice before we add or replace a sub-processor that processes customer personal data.

Sub-processor	Purpose	Data categories	Region
---------------	---------	-----------------	--------

Supabase (Lovable Cloud backend)	Primary database, authentication, file storage, edge functions	All customer data at rest and in transit	Configurable; default US
Google (Gemini via Lovable AI Gateway)	LLM inference for the coach, summaries and artifact generation	Prompts and the context sent per request (see section 5)	Google Cloud regions
Recall.ai	Optional meeting bot joining Zoom, Teams or Meet	Audio, video, transcript, participant metadata — only when a user starts a bot	US
Stripe	Subscription billing and payment method storage	Billing contact, card token, invoice history; no card number touches Quorval	US / EU
Resend / email provider	Transactional email: invites, notifications, magic links	Recipient email and message body	US / EU
Edge/CDN via Lovable hosting	Static asset delivery and TLS termination	Public assets and request metadata	Global edge

Residency note. Because Quorval runs on shared managed infrastructure, it inherits the residency, retention and compliance posture of those providers for the services it consumes. Regimes that require personal data to remain in a specific jurisdiction, or to avoid US-controlled providers, should discuss requirements with us before onboarding; a dedicated-instance offering is on the roadmap.

3. Tenancy, isolation and access control

Quorval runs a shared multi-tenant PostgreSQL database. Isolation is enforced at the row level by policies scoped to `organization_id`, plus a pattern that requires the caller's authenticated user id to appear in the target organization's membership before any read or write is permitted. Every public-schema table ships with row-level security enabled and explicit grants.

Role hierarchy is stored in a dedicated roles table — never on the user or profile record — and checked through a security-definer function so that policy evaluation cannot be used for privilege escalation.

Cross-tenant access is possible only in two audited, opt-in scenarios: a designated vendor-support user impersonating a specific organization for a support ticket, gated by an organization-level toggle and written to an immutable audit log; and a stakeholder or external assignee reaching a single project through a signed, revocable link.

New in version 1.1 — isolation is now tested, not asserted. A deterministic database probe inspects every tenant-scoped table for enabled row-level security, policy coverage, absence of blanket-read policies, and organization scoping inside shared retrieval functions. It returns a binary factual result rather than a model's opinion. The most recent run passed 212 of 212 checks. The probe runs on a monthly schedule and its results are retained as evidence.

Scope note. The probe suite evidences that the access-control configuration behaves as designed. It is complementary to, not a substitute for, adversarial penetration testing, which is scoped and planned. Customers with regulated workloads — HIPAA-covered PHI, GxP or EU public-sector data — should perform their own risk assessment before uploading regulated data.

4. Storage, encryption and key management

In transit: TLS 1.2 or better is enforced end-to-end for customer traffic, backend calls, edge functions and sub-processor calls.

At rest: managed disk encryption on the hosted PostgreSQL cluster and object storage using AES-256 with provider-managed keys. Selected sensitive columns, such as stored secrets and portal access tokens, are additionally encrypted at the application layer.

New in version 1.1 — file storage authorization. Every bucket holding customer content is private and read and write access is scoped to the owning project or organization, checked by a security-definer function rather than by a blanket authenticated-user rule. Branding assets are stored under an organization-keyed path so one tenant cannot overwrite another's files.

Anonymous read access to customer content buckets has been removed.

Keys and secrets: application secrets live in the platform secret store, never in source control or the browser bundle. The database service-role credential is held by the platform and is not available to Quorval operators or to customers on the shared-tenant plan.

Key management scope. Data-at-rest keys are managed by the underlying infrastructure provider. Customer-managed encryption keys (CMEK) are not offered today; customers with a CMEK mandate should raise it during review.

5. LLM data use — the part auditors care about most

Quorval uses large language models, currently the Google Gemini family accessed through the Lovable AI Gateway, to power the AI coach, meeting summarization, artifact drafts, risk suggestions and change-management communications.

What we send to the model, per request: only the context needed to answer the prompt — typically the user's message, the current project's relevant records, and short system instructions. We do not stream your entire database, another organization's data, or unrelated projects into the model.

What we do not do:

- We do not use your prompts, uploads, transcripts or generated artifacts to train, fine-tune or improve any third-party foundation model.
- We do not use your data to train a Quorval-owned model. Quorval does not operate its own foundation model.
- We do not sell customer data or share it with advertisers or data brokers. There is no advertising pixel on authenticated product pages.
- We do not enable model providers' opt-in training programmes on your behalf. Requests go to the provider's paid API tier, where inputs and outputs are not used for training under the provider's enterprise terms.

Retention at the model provider. The paid Gemini API has a documented no-training-retention posture and a short operational logging window for abuse monitoring. That window is set by

the provider, not by Quorval, and can change; we track their terms and will notify customers of material changes.

Prompt and response storage in Quorval. Coach conversations, generated summaries and suggested artifacts are stored in the customer's workspace like any other content, scoped by row-level security, and inherit the customer's retention policy. Deleting a project or record deletes its associated AI-generated artifacts.

Meeting transcripts. When a user sends a recording bot to a meeting, audio is captured by Recall.ai, transcribed, and delivered to Quorval through a webhook whose signature we verify with HMAC-SHA256. Recording files are stored in the customer's private storage bucket. Recall.ai's own retention and residency terms apply to copies held on their side.

Scope of LLM use. Model outputs can be wrong; AI-generated content is labelled and human review is required before it becomes an approved artifact or an e-signed record. Assurance over third-party model handling of prompts rests on contractual terms with the provider rather than on technical enforcement by Quorval. Text a user enters into the coach is transmitted to the provider's API, so customers should train users accordingly and may disable AI features where required.

6. Authentication, session and audit

Authentication supports email and password, Google, Apple, and SAML single sign-on for enterprise customers. Passwords are hashed by the identity provider using industry-standard key derivation; Quorval never sees plaintext passwords. Password creation is screened against known-breached credential lists.

New in version 1.1 — multi-factor authentication. Time-based one-time-password enrolment is available to all users and is enforced for privileged roles: organization owners, organization administrators and vendor support staff cannot reach privileged surfaces without completing a second factor. Recovery codes are single-use and their redemption is logged and alerted.

Audit trail. Authentication events, privileged actions, role changes, impersonation sessions, document generation, exports and signature events are written to append-only logs with actor, target, timestamp, IP address and user agent, and are retained on a defined schedule — 24 months for authentication, authorization and sensitive data access. Application code has no path that deletes or rewrites an audit record.

7. Operational monitoring of our own automations (new in 1.1)

Scheduled work in Quorval — reminder delivery, calendar imports, retention enforcement, report delivery, evaluation runs — no longer calls its target function directly. Each invocation passes through a dispatcher that records a start row, runs the target, and closes the row with status, duration, items processed and any error. Ten recurring jobs are registered with an expected interval and criticality.

A health view derives a verdict of healthy, stale, hung, failing or never-run for each job by comparing its last successful run against its expected interval. A watchdog evaluates that view

every fifteen minutes and notifies administrators when a job is stale, hung or repeatedly failing. The practical effect for customers is that an automation which stops firing is detected from data within minutes rather than from a support ticket days later.

A control register holds fifteen recurring security and AI-governance controls with an owner, a frequency, the evidence required and a verification mode. A daily scheduler creates each period's activity, automatically verifies the controls a machine can check, and flags overdue items to their owner. Manual controls are visible as outstanding until a person performs them.

Scope note. Automated verification evidences that the control automation executed as configured; correctness of the control itself is established by review and, where applicable, by test. The register was created in August 2026, and the multi-period history required for a SOC 2 Type II observation window is accruing from that date.

8. How we evaluate our own AI (new in 1.1)

Quorval applies its own AI governance model to its own AI systems. Eight AI features are registered with a stated purpose, risk classification and human oversight point. The exact production prompt text is held in a versioned registry, and evaluation suites are bound to the prompt and model each feature actually uses in production, so an evaluation exercises the deployed system rather than a stand-in.

Evaluation suite	Mode	Threshold	Latest result
Tenant isolation in retrieval	Deterministic	100%	Pass — 212 of 212 probes
Coach methodology grounding	Model-judged	90%	Fail — 55%; open finding, prompt remediation in progress
Retrieval citation integrity	Model-judged	90%	No clean recorded run; no baseline yet
Meeting extraction precision	Model-judged	85%	No clean recorded run; no baseline yet

Results are published in full, including suites below threshold. The isolation result is strong evidence because it inspects access-control state directly and is reproducible. The model-judged results are weaker by nature — a judge model can be wrong in either direction — and are used as a monitoring signal and regression tripwire, not as an assurance statement. A suite scoring below threshold opens an internal AI incident. No evaluation result is presented to customers as a certification, and none is used to claim that our AI systems are accurate.

9. Retention, deletion and data-subject rights

Retention policies are configured per organization and enforced by a daily job that purges expired records; the job's execution is logged like every other automation. Deleting a project deletes its associated records, generated artifacts and uploaded files. Account deletion requests are honoured within 30 days, subject to legal-hold and billing-record obligations.

GDPR and UK GDPR. As processor, Quorval assists the controller with access, rectification, erasure, restriction, portability and objection requests. As controller of its own account data, Quorval handles requests directly through the contact in section 12.

APPI (Japan). Quorval supports disclosure, correction and cessation of use for retained personal information. Cross-border transfers to the United States are disclosed in section 10 and form the basis on which a controller can obtain consent under Art. 28 where required.

CCPA and CPRA. Quorval does not sell or share personal information for cross-context behavioural advertising. Consumers exercise access and deletion rights through the controller.

10. International transfers

The default hosting region for the managed database is the United States. Data is transferred to the United States and to the Google Cloud regions serving the Gemini API. For EU and UK customers, Quorval relies on the Standard Contractual Clauses (2021/914) and the UK Addendum, incorporated by reference in our data processing agreement, together with the measures described in sections 3, 4 and 5. For Japanese customers subject to APPI cross-border rules, the destination countries above are disclosed and the specific safeguards required for Art. 28 consent-based transfers are provided on request.

Data residency. Quorval does not currently offer EU-only or Japan-only residency on the shared-tenant plan. Customers with in-country processing requirements should discuss options with us before onboarding.

11. Certifications and current posture

The table below states our current position against each framework. Only certifications actually held or formally underway are listed.

Framework	Current status	Notes
SOC 2 Type I	In preparation; not audited	Policy set issued, control register operating, technical controls implemented. Remaining work is manual operating evidence: first access review, restore test, incident tabletop, penetration test.
SOC 2 Type II	Not audited	Requires an observation window that began accruing in August 2026. Underlying infrastructure providers hold SOC 2.
ISO/IEC 42001 (AI management)	In preparation; not certified	AI inventory, prompt registry, human-oversight checkpoints and scheduled evaluations are live. Signed AI policy, impact assessments, internal audit and management review are outstanding.
ISO/IEC 27001	Not certified	We follow aligned practices; certification not undertaken.
ISO 9001:2015	Aligned, not certified	Document control, traceability and corrective-action patterns are built into the product.
HIPAA	No BAA offered today	Do not upload PHI. A HIPAA-aligned tier is on the roadmap.
GDPR / UK GDPR	Processor obligations met via	No supervisory-authority findings against

	DPA, SCCs and this statement	Quorval.
APPI (Japan)	Processor obligations met via DPA and the disclosures in section 10	No PPC findings against Quorval.
21 CFR Part 11	Product features aligned (e-signature, audit trail)	Customer validation (IQ/OQ/PQ) remains the customer's responsibility.
FedRAMP / IL-x	Not authorized	Do not use for US federal controlled data.
PCI DSS	Out of scope	Stripe handles card data; Quorval stores no cardholder data.

12. Incident response, vulnerability reporting and contact

Quorval maintains an incident response process with severity classification and post-incident review. For a personal-data breach affecting customer data, Quorval notifies the affected controller without undue delay and within 72 hours of confirmation, providing data categories, approximate record counts, likely consequences and mitigation steps — sufficient for the controller to meet its own GDPR Art. 33 or APPI Art. 26 obligations.

Detection note. Because Quorval runs on managed sub-processors, some breach detection depends on their disclosures to us. The 72-hour clock starts when Quorval confirms the incident, not when a sub-processor first detects it. A documented tabletop exercise is scheduled as a registered control.

Operating entity: CreativePrinTech LLC (dba Quorval), United States. Security and privacy contact: privacy@quorval.com. Vulnerability reports: privacy@quorval.com with “Vulnerability report” in the subject line; we acknowledge within two business days. Requests handled include DPA execution, sub-processor updates, security questionnaires (SIG Lite, CAIQ), and data-subject request assistance. An EU or UK Art. 27 representative is available on request; an APPI-designated representative in Japan is not currently appointed and should be raised before contracting.

13. What to do with this document

This statement is intended to be attached to a vendor security review, a record of processing activities entry, or a data protection impact assessment. It is not a substitute for the Quorval data processing agreement, which contains the binding commitments. Where this document and the DPA conflict, the DPA governs. Where this document and reality conflict, reality governs and we owe you a correction — please tell us.