

Hydra Security Overview

Radiant InnovaTech · Version 1.1 · 22 October 2026 · For customers and their security teams

■ 1. What this document is

This document describes how Hydra is built and how Radiant protects it and the data customers put into it. We mark every control with one of three statuses so that you can see exactly where we are: **In place** (running today), **In progress** (being built or confirmed now) and **Planned** (scheduled, not yet started). Where a control is not yet in place we say so. We update this document when things change and can share the detailed answers to a standard security questionnaire on request.

■ 2. Summary

Area	Status	In one line
Advisory-only operation	In place	Hydra recommends and a person decides. It does not change a customer's systems.
Sign-in and access control	In place	Managed sign-in with multi-factor authentication available; five roles with least privilege.
Separation between customers	In progress	Every record belongs to a workspace and every query is scoped to it; database-level enforcement is being confirmed.
Encryption	In place	HTTPS only; data encrypted at rest by the cloud provider; stored integration credentials encrypted again by Hydra.
Audit trail	In place	Hydra records its actions and its users' decisions in a tamper-evident trail.
AI safeguards	In place	Verdicts cite evidence, alert content is treated as untrusted data, no training on customer data.
Backups	In place	Automated daily backups, a manual backup before each release.
Recovery and availability	In progress	Single-zone database; high availability, point-in-time recovery and recovery drills are planned.
Incident response	In progress	Written procedure and yearly test being put in place.
Independent testing	Planned	Third-party penetration test scheduled.

Area	Status	In one line
Certifications and licensing	In progress	Licence from the Cyber Security Authority being prepared; ISO 27001 certification planned; SOC 2 scoping planned.
Data protection registration	Planned	Radiant is not yet registered with the Data Protection Commission and will register once the company is registered, before any customer data reaches Hydra.

■ 3. How Hydra is built

Hydra has four parts: a web application that analysts use; an application programming interface (API) that holds the business rules and talks to the customer's security tools; an AI investigation service that reads an alert, gathers evidence through the API's read-only tools and proposes a verdict; and a database. They run on Google Cloud in a managed Kubernetes cluster, with the database on Google's managed PostgreSQL service. Customers reach Hydra over HTTPS at our own domain. By default Hydra is hosted in the United States (Iowa); hosting in Ghana is planned.

■ 4. Advisory-only operation and the AI (In place)

- Hydra is advisory-only. It does not isolate hosts, disable accounts, block addresses or change anything in the customer's systems. Any action it suggests goes to a person for approval, and the person carries it out in the customer's own tools.
- An optional triage feature that can close routine alerts is off by default for evaluations and runs in a record-only "shadow" mode first. It never closes an alert rated critical, has an hourly limit, and its decisions can be reversed.
- Every verdict must cite the evidence it relies on. A fixed layer of detection rules checks the model's answer, and Hydra answers "needs context" when it cannot decide.
- Alert content is treated as untrusted data: the AI is instructed not to follow any instruction found inside an alert, which protects against prompt injection.
- We do not train AI models on customer data and we do not let our model providers do so. Customers choose whether alerts are read by a hosted model from Google or by an open model that we run ourselves.
- Analysts can mark each verdict Agree or Disagree. The record of agreement and of any attack that Hydra dismissed is reported to the customer weekly during an evaluation.

■ 5. Identity and access (In place)

- Sign-in is handled by an identity provider (Clerk), with multi-factor authentication available and enforceable for a whole workspace.
- Hydra has five roles: viewer, analyst, manager, admin and owner. Reading, investigating, approving actions and changing configuration each need a different role, so a person gets only what they need.

- The AI service uses a separate service identity with read-only access. It cannot approve or carry out anything.
- Radiant staff access to production is limited to named people who need it. We are putting a formal access review in place (In progress).

■ 6. Protecting data

- **Separation between customers (In progress).** Each customer's data sits in its own workspace and every query is scoped to that workspace in the application. Database row-level security policies are applied in production; we are confirming that they are enforced under the application's database account, and will report the result.
- **Encryption (In place).** Traffic to Hydra uses HTTPS only, with a managed certificate and a strict transport security header. The database and storage encrypt data at rest. Credentials that customers store in Hydra to connect their tools are additionally encrypted by Hydra with AES-256-GCM.
- **Secrets (In place).** Keys and passwords are kept out of the source code and delivered to the services at run time.
- **What we ask for (In place).** Hydra reads alert details, not file contents, email contents, passwords or payment card data, and we ask customers not to send those. Connectors read from the customer's tools with read-only access.
- **Return and deletion (In place by agreement).** When an agreement ends we delete the customer's data within 30 days and from backups within 90 days, and confirm in writing if asked.

■ 7. Infrastructure and network

- Services run in containers that run as non-root users, with a restricted set of system privileges (In place).
- A default-deny network policy blocks traffic between services unless it is explicitly allowed, so a compromise of one service does not give free movement to the others (In place).
- Calls from Hydra to a customer-configured address are checked against a rule that blocks private networks and the cloud metadata service unless the deployment is configured for them, which prevents server-side request forgery (In place).
- Standard security headers are set on the web application: HSTS, a content security policy, frame denial, content-type protection and a strict referrer policy (In place). We are tightening the content security policy further (In progress).
- A web application firewall and shared rate limiting at the network edge are planned. Rate limiting is applied in the application today (Planned for the edge).

■ 8. Building and changing the software

- Changes are made by pull request and tested before release (In place). Review by a second person on every change is being made mandatory (In progress). The code base has more than 300 automated tests, including tests of tenant separation and of the safeguards that stop Hydra acting on its own.
- Releases use images that are identified by a commit and cannot change after they are built, and a release script that rolls back by itself if the new version does not start (In place).
- A database backup is taken before each release (In place). Database changes are applied automatically as the new version starts.
- Dependencies are audited on each release, and high or critical findings are fixed before the release or within 14 days (In progress: we are completing an update of the main framework and database packages now).
- Automated checks on every pull request and protected branches are being set up (In progress).

■ 9. Logging, audit and monitoring

- **Audit trail (In place).** Every action Hydra takes, every decision a user makes and every configuration change is recorded in a tamper-evident chain, so that an altered or deleted record can be detected. Customers can see their workspace's audit trail.
- **System logs (In place).** Application and platform logs go to Google Cloud logging.
- **Alerts (In progress).** We are adding alerts for crashes, failed sign-ins and unusual traffic, and a rota for who responds.

■ 10. Backups, recovery and availability

- Automated daily database backups are kept for 14 days, and a manual backup is taken before every release (In place).
- A restore from backup into a clean database is tested at least twice a year. The first drill is scheduled (In progress).
- Point-in-time recovery, a database that survives the loss of a zone, more than one running copy of the service, and agreed recovery targets are planned before we offer a service level agreement (Planned). During an evaluation we do not promise a service level.

■ 11. Security incidents

Radiant is writing an incident response procedure that covers how we detect, contain and investigate an incident, who decides, and how we notify customers and regulators (In progress). The commitment we make now, in our data processing agreement, is to tell a customer without undue delay and within 48 hours after we become aware of a breach affecting its data, with the details we have, and to help the customer to investigate and report it.

■ 12. Independent assurance

- Independent penetration test: planned during the first pilots (Planned).
- ISO 27001 and SOC 2: we are scoping both, and have not yet started (Planned).
- Licence: Radiant plans to run Hydra for customers itself, under a licence from Ghana's Cyber Security Authority for managed cybersecurity services. We are preparing the application, and company registration is in progress (In progress). Radiant will not provide Hydra to a customer until it is licensed or the Authority has confirmed in writing that the activity does not need a licence. Today Radiant does not operate a customer's security operations centre: the customer's own analysts run the monitoring and Hydra advises.
- Data protection registration: Radiant is not yet registered as a data controller and processor with Ghana's Data Protection Commission. We will register once the company is registered, and before any customer's real data reaches Hydra (Planned).

■ 13. Where data goes and who handles it

Provider	What it does	Location
Google Cloud	Hosts Hydra, its database, storage and backups	United States (Iowa) by default
Google Vertex AI	Hosted AI model, only if the customer chooses it	United States
Clerk	Sign-in and accounts	United States
Brevo	Email notifications, only if enabled	European Union
mNotify	Text-message notifications, only if enabled	Ghana

We will tell customers 14 days before adding a provider that handles their data.

■ 14. Reporting a security issue and contact

To report a vulnerability or ask a security question, write to security@radiantinnovatech.com. We acknowledge reports within two Business Days and do not take action against people who report in good faith. For privacy matters, write to privacy@radiantinnovatech.com.