OPERATOR MANUAL · v1.0

Global Rail Cyber Security Console

Real-world incidents, translated into checks, controls, and action.

AUDIENCE · security leads · engineers · auditors · superadmins
SCOPE · console operation, day-to-day workflows, admin & integrations
REVISION · June 2026

00 · CONTENTS

How this manual is organized

Read straight through for a complete tour, or jump to the section you need. Every chapter mirrors a section of the in-product docs at */docs/console*, so the manual and the live guide stay in lockstep.

01 · OVERVIEW

Run the Console with confidence

The Global Rail Cyber Security Console turns real-world incidents into checks, controls, and action — without ever pushing automated changes into your operational environment.

What it does

The Console gives a railway-grade security team one workspace for four jobs:

- **SecOps** — scoped, safe assessments against your own assets, with rail-safety guardrails.
- **Live SOC** — triage incoming detections, acknowledge, escalate, and tie them to incidents.
- **AI Audit** — map AI systems to TS 50701, NIS2, ISO 27001, NIST CSF, and produce evidence.
- **Threat Library** — a curated catalog of rail-relevant threats, convertible into checks.

Who it's for

Security leads prove posture, plan remediation, and brief leadership. **Engineers** investigate, triage, and fix. **Auditors** pull evidence packs. **Superadmins** manage tenants, identity, and engine health.

Human-in-the-loop guarantee
The Console never pushes automated changes to customer environments. Every detection is reviewed by a human before it ships. Every remediation leaves the building under your own change-control process.

How this manual maps to the product

Each chapter matches a section of the in-product docs at **/docs/console**. When something changes in the product, the docs site is the source of truth — treat this PDF as the offline companion for operators who need a single artefact.

02 · GET STARTED

Your first 15 minutes

Sign in, pick a tenant, connect your first asset, and run your first assessment. This chapter is the offline equivalent of the in-product Quickstart.

1. Sign in

Open the Console URL provided by your administrator. Authenticate with email + MFA or your organization's SSO provider. First-time superadmins receive a one-time invite link.

2. Pick a tenant

The tenant switcher lives in the top-left of the side navigation. Superadmins see *All tenants* plus every workspace they belong to. All data — assessments, findings, evidence — is scoped to the active tenant.

Demo workspace
Every account has read-only access to the **Acme Rail Demo** tenant. Use it to explore the product without touching production data.

3. Connect an asset

Navigate to **Assets** → **Add asset**. Choose a cloud connector, install an on-prem agent, or enter the asset manually. CMDB sync runs nightly once configured.

4. Run your first assessment

From **SecOps** → **Run new**, pick a scope, a profile, and a schedule. Read-only checks run by default; intrusive checks require an explicit change window.

Roles at a glance

Role	Can do
Viewer	Read findings, dashboards, evidence.
Operator	Run assessments, triage Live SOC, open tickets.
Admin	Manage users, integrations, SSO, allowlists.
Superadmin	All of the above across every tenant; engine + competition tools.

03 · WORKFLOWS

What good looks like, week to week

Choosing what to test

Scope deliberately. Test public-facing assets, recently changed systems, and assets that touch passenger or safety data. Avoid testing safety-critical control systems without an explicit change window and a rollback plan.

Safe testing on operational systems

The Console enforces three guardrails on rail environments:

- **Read-only mode** by default — no writes, no exploitation, no DoS.
- **Allow-listed checks** — only checks marked safe-for-OT run inside a change window.
- **Rollback hooks** — every intrusive check ships with a documented rollback path.

Interpreting results

Severity is not CVSS. The Console combines exploitability, rail-safety impact, and observed prevalence into a single 0-100 score. See the Reference chapter for the formula.

Closing the loop

A finding's lifecycle: *open* → *triaged* → *assigned* → *fixed* → *verified* → *closed*. Every transition is logged with the actor, timestamp, and (optional) evidence link. Re-test on demand or on a schedule to confirm the fix held.

Rail safety note

Never close a finding on a safety-critical asset without a signed-off re-test. The Console will block the transition if the asset is tagged *safety-critical* and the re-test is missing.

04 · HOW-TO GUIDES

Task-focused recipes

Connect assets

Cloud connectors (AWS, Azure, GCP) pull asset inventory via read-only IAM roles. On-prem agents register over outbound TLS — no inbound ports required. Manual records and CSV import cover anything the connectors miss.

Set up authentication

Email + password is enabled by default. Add MFA in **Settings** → **Security**. For SSO, connect OIDC or SAML in **Admin** → **SSO**. Step-up auth is required for destructive actions (delete tenant, rotate API keys, change SSO config).

Run an assessment

Pick a target scope, choose a profile (Baseline / Deep / Custom), set a schedule, and launch. Progress streams live; cancel at any time. The output is a set of findings, an evidence pack, and a machine-readable JSON export.

Use the Threat Library

Browse the curated catalogue under **Threat Library**. Each threat carries a rail-context summary, indicators, and a one-click conversion into a check that runs in your next assessment.

Operate Live SOC

Detections arrive in the inbox. Acknowledge to take ownership; escalate to open an incident; link to an existing incident to merge context. The SOC view shows source, rule, asset, and the last 24h of related activity.

Run an AI audit

Register an AI system, pick the applicable frameworks, and walk the controls. The Console drafts evidence prompts; you attach artefacts. The final report is a downloadable PDF + a JSON export for your GRC tool.

Use the AI Agent

Open the assistant from the bottom-right of any screen. Ask in natural language: *“summarize today's critical findings”*, *“draft a status email for the SOC lead”*, *“take me to the Wiz integration”*. The agent reads from the docs and from your visible context only — never across tenants.

Export evidence

From any finding, assessment, or audit, choose **Export** → PDF, CSV, or signed JSON. Audit packs include a manifest with SHA-256 hashes for every file.

05 · REFERENCE

Reference

Severity model

Final severity is computed as:

```
severity = clamp(
0.45 * exploitability
+ 0.35 * rail_safety_impact
+ 0.20 * prevalence,
0, 100)
```

Bands: **0–24 Info**, **25–49 Low**, **50–74 Medium**, **75–89 High**, **90–100 Critical**.

Finding schema (abridged)

```
{
"id": "uuid",
"tenant_id": "uuid",
"asset_id": "uuid",
"severity": 0-100,
"state": "open|triaged|assigned|fixed|verified|closed",
"framework_refs": ["TS-50701:8.3.2", "NIS2:Art.21"],
"evidence": [{ "sha256": "...", "path": "..." }],
"created_at": "ISO-8601",
"updated_at": "ISO-8601"
}
```

Frameworks

- **EN/TS 50701** — rail cybersecurity baseline.
- **NIS2** — EU directive for essential entities.
- **ISO/IEC 27001:2022** — ISMS controls.
- **NIST CSF 2.0** — functions and categories.

Attack coverage

The Console covers: external attack surface, identity & access, cloud misconfiguration, credential hygiene, third-party exposure, AI-system misuse, and rail-specific control-plane patterns. It does **not** perform exploitation, lateral movement, or social engineering.

Roles and permissions

See the table in Chapter 02. Permissions are tenant-scoped; superadmins are cross-tenant.

API

Read-only REST API at **https://api.cybersecurity.globalrailsuite.com/v1**. Bearer-token authentication, 600 req/min per token, cursor pagination.

```
GET /v1/findings?tenant_id=...&state;=open&limit;=100
Authorization: Bearer <api-token>
```

06 · ADMIN

Admin

Users

Invite by email; promote, demote, or remove from **Admin → Users**. Active sessions are listed per user with the ability to force sign-out. Every change is captured in the audit log.

SSO

Connect OIDC or SAML providers. Set the default role for newly provisioned users. Group-to-role mapping is supported for OIDC *groups* and SAML *memberOf*.

IP allowlist

Restrict Console access to specific egress IP ranges. Allowlists apply per tenant and are enforced at the edge. Break-glass access is available to superadmins from a registered out-of-band device.

Audit log

Tamper-evident, append-only, hash-chained. Exportable as signed JSON. Retention is 365 days by default; extend per tenant under **Admin → Data residency**.

Data residency

Pick EU, UK, or US storage at tenant creation. Residency cannot be changed later without a documented migration. All processing happens in the same region as storage.

07 · INTEGRATIONS

Integrations

Slack

Route findings, SOC events, and audit milestones into Slack channels. Per-tenant routing rules; supports rich blocks with deep links back into the Console.

Jira

Open and track remediation tickets from a finding. Bi-directional sync keeps status, assignee, and comments aligned. Project mapping is per-tenant.

SIEM

Forward Console events to Splunk, Sentinel, or any syslog/CEF-compatible SIEM. Choose JSON-over-HTTPS or syslog. Field mapping is documented in the SIEM page in-product.

Wiz

Ingest Wiz findings into the Console and unify them with rail-specific context. Severity is re-scored using the Console's model so the queue stays consistent.

Webhooks

Receive Console events at your own HTTPS endpoint. Payloads are HMAC-SHA256 signed; verify with the per-webhook secret. Retries: 5 attempts with exponential backoff.

08 · SUPERADMIN

Superadmin tools

Cross-tenant capabilities reserved for the platform's own operators. Use sparingly — every action is logged and visible to the customer.

Tenant management

From **Clients**, create, suspend, or merge tenants. Provision the initial admin by email; the customer takes over from there. Suspension preserves data and freezes all writes.

Economics

Superadmin → **Economics** shows assessment volume, AI token cost, and unit economics per tenant. Use it to size pricing changes and to spot runaway consumption.

HKSECAI — competition module

HKSECAI is a one-off feature for the Hong Kong Security AI competition (Jul 4, 2026). It is not a customer-facing capability and lives entirely under **Superadmin** → **HKSECAI**. The module:

- Accepts a competition **brief** (paste, file, or URL) and extracts targets via OpenClaw.
- Stages extracted targets into a sanctioned engagement on a chosen tenant.
- Runs a signed-HMAC **Discover / Solve / Submit** loop against the HKSECAI engine.
- Supports two engine modes: **Single** (proven, default) and **Swarm** (Opus 4.8).
- Never auto-submits flags — every submission is human-approved.

Engine connectivity

Use **Ping engine** on the HKSECAI page. The button runs an unsigned *GET /healthz* (reachability) followed by a signed *POST /hksecai/ping* (HMAC + allowlist). If healthz fails, the tunnel or hostname is down; if ping fails, the issue is the HMAC secret or the allowlist.

Swarm mode budget knobs

```
engine: "single" | "swarm"
workers: 1-16 (swarm only)
worker_timeout: 30-1800s (swarm only)
max_rounds: 1-20 (swarm only)
wall_timeout: seconds (always applies)
```

Scope lock

The engine refuses any target not on the allowlist. The allowlist is populated only from a staged competition target. There is no path from a customer asset to the HKSECAI engine.

09 · TROUBLESHOOTING

When something looks wrong

I can't sign in

Confirm the URL, then your MFA device clock. If SSO, check with your IdP admin that the user is in the provisioned group. Superadmins can issue a break-glass link.

Assessment is stuck in *running*

Open the run detail; the step log shows the last heartbeat. Steps with no heartbeat for > 5 minutes are auto-cancelled. If the asset is unreachable, the connector page shows the underlying error.

A finding looks wrong

Open the finding, click **Evidence**. Every check ships with the raw signal that produced the finding. If the signal is correct but the verdict is wrong, mark it *false positive* with a reason — the model learns from this.

Webhook isn't firing

Check **Integrations** → **Webhooks** → **Deliveries**. Failed deliveries show response code and body. Retries: 5 attempts, exponential backoff up to 1 hour.

HKSECAI engine ping fails

Run the **Ping engine** button. If *healthz* fails, the engine host is down (escalate to platform). If *healthz* passes but *ping* fails, the HMAC secret is mismatched or the caller IP is not on the allowlist.

10 · SUPPORT

Getting help

In-product: open the assistant (bottom-right) and ask. The assistant can deep-link to the right screen and draft a support email on your behalf.

By email: **support@cybersecurity.globalrailsuite.com**. Include tenant ID, the URL of the screen you were on, the action you took, and what you expected.

For incidents that affect production safety, follow your organization's internal incident process first. The Console is an analytical tool — it does not replace operational incident command.