PORT CYBERDEFENSE

CONFIDENTIAL • TECHNICAL BRIEF

GHOST RED

v2.5.0 Platform

Autonomous AI Adversary — Architecture, Integration & Roadmap

PRODUCT	STATUS	DOCUMENT	CLASSIFICATION
GHOST RED v2.5.0	**In Production**	**Technical Brief**	**Confidential**

Prepared for: Security Architects • CISOs • CTOs • Partners

Port Cyber Defense • Málaga, Spain • Strasbourg, France
info@port-cyber-defense.com • port-cyber-defense.com

Table of Contents

01 Executive Summary

What GHOST RED is and why it exists

GHOST RED is an autonomous AI adversary platform. It is not a vulnerability scanner, not a penetration testing tool, and not a consultancy deliverable. It is a continuous, production-grade SaaS system that simulates nation-state-level cyberattacks against enterprise environments — chaining vulnerabilities into real kill paths, mapping every finding to MITRE ATT&CK, and auto-generating compliance evidence for DORA, NIS2, ISO 27001, and PCI DSS.

Traditional red team engagements cost €30,000 to €100,000 per engagement, take weeks to complete, and produce a snapshot report that is outdated the moment it is delivered. GHOST RED replaces that model with continuous autonomous testing at a monthly SaaS price point, with compliance reports regenerated on every scan.

Platform Status (v2.5.0)

v2.5.0	6/6	1,250+	24/7
CURRENT RELEASE	SUBSYSTEMS HEALTHY	CVES SYNCED DAILY	AUTONOMOUS SCANNING

What This Brief Covers

- Platform architecture — the Celery + Redis + Claude AI + NVD/KEV stack, two distinct worker pools
- The 9 scan phases from reconnaissance through impact, with technical detail per phase
- The AI Chain Engine — how CVEs are chained into executable attack paths
- Compliance framework coverage — 4 live today, 5 on the 2026 roadmap
- Deployment models — SaaS, Sovereign on-premise, and air-gapped options
- Integration roadmap — SIEM, ITSM, identity, and vulnerability scanner connectors
- Quarterly product roadmap through end of 2026
- Security and data handling practices

WHO THIS BRIEF IS FOR

This document is intended for security architects evaluating GHOST RED for enterprise deployment, CISOs considering adoption, CTOs reviewing integration feasibility, and partners assessing joint go-to-market potential. It assumes technical familiarity with offensive security concepts and enterprise security operations.

02 Platform Architecture

The technical stack under the hood

GHOST RED is built as a distributed, event-driven system. The architecture is optimised for three things: continuous autonomous operation, horizontal scalability of scan workers, and deterministic compliance reporting.

High-Level Architecture

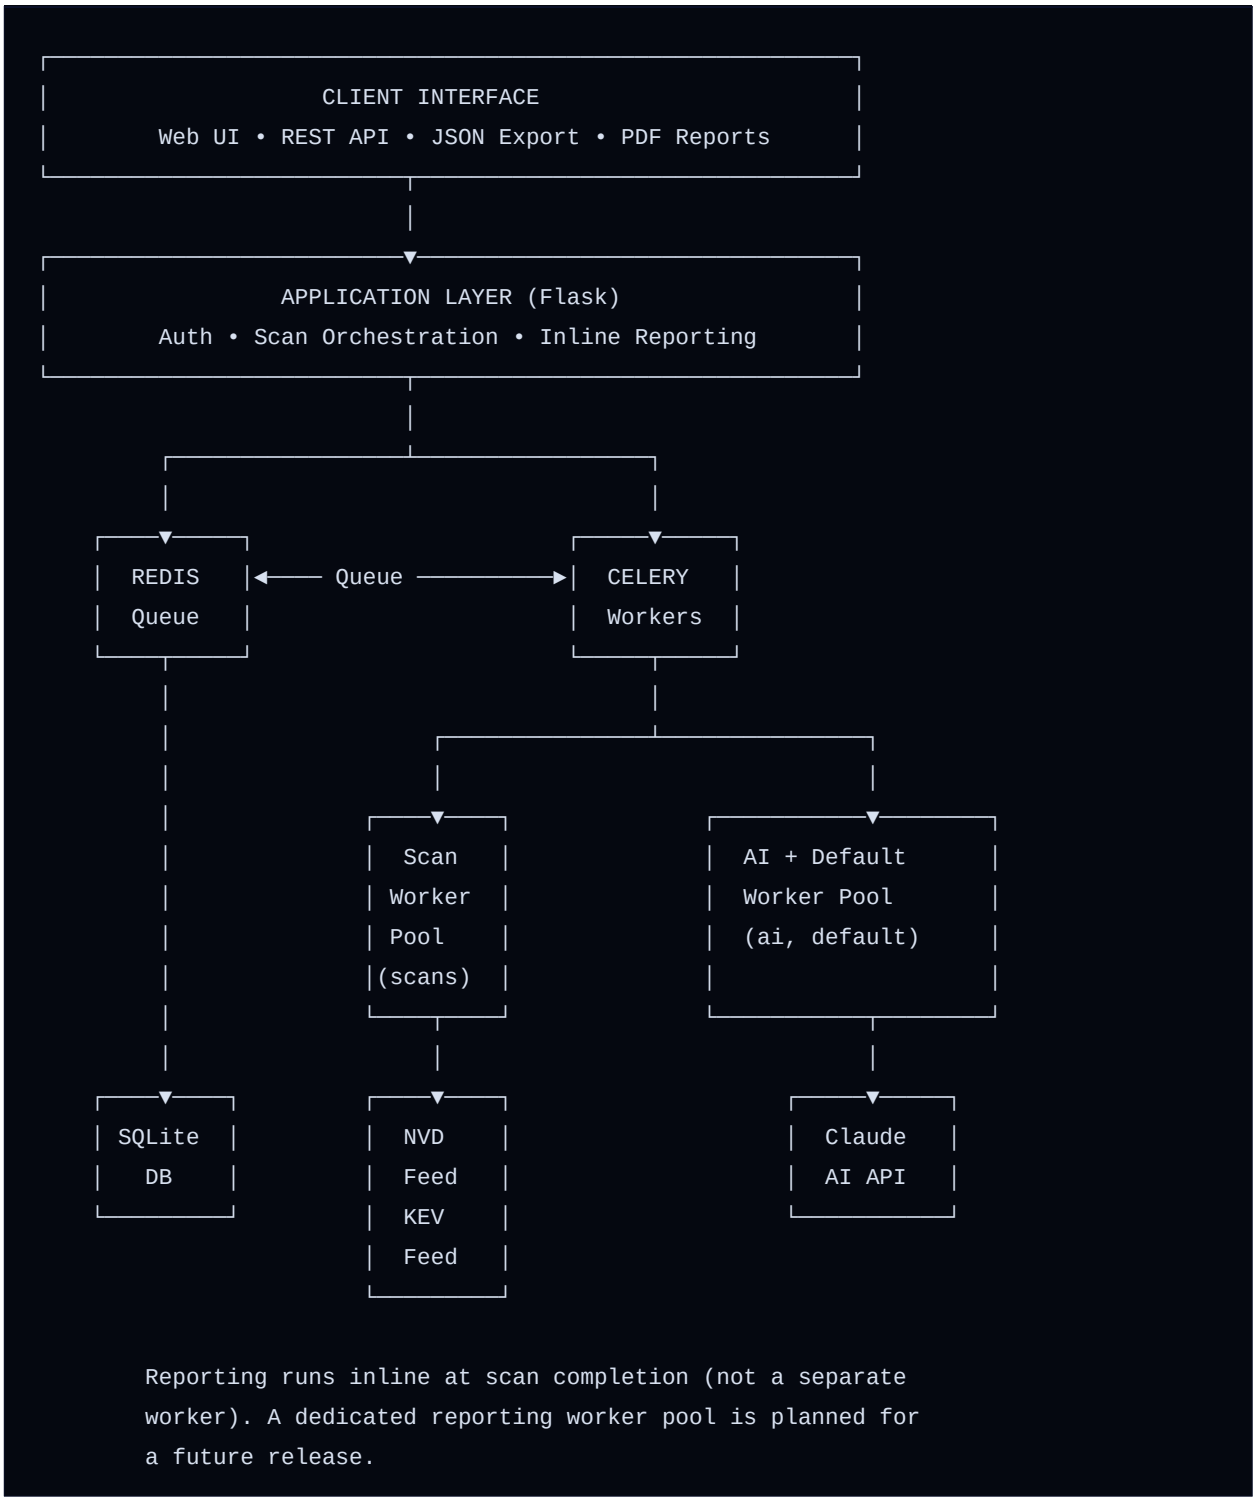

Core Components

Application Layer (Flask)

A Python Flask application exposes the web UI, REST API endpoints, and scan orchestration logic. This layer handles authentication, scan scheduling, and assembly of the final compliance reports. The application is stateless and can be horizontally scaled behind a load balancer.

Redis Queue

Redis serves two roles: as the message broker for Celery workers, and as a cache for short-lived scan state. Scan requests are placed on the Redis queue where they are consumed by worker processes running on separate machines (or containers). This decouples scan execution from the user-facing API.

Celery Workers

Celery runs two worker pools today, each subscribed to its own queue set:

WORKER POOL	RESPONSIBILITY
Scan Worker Pool (queue: scans)	Executes the 9 scan phases against target infrastructure. Runs port scans, service enumeration, vulnerability probes, and evidence collection. Dedicated to scan execution so heavy scans cannot block AI reasoning.
AI + Default Worker Pool (queues: ai, default)	Generates attack chains by consulting the Claude AI API. Takes discovered CVEs and network context as input, outputs ranked multi-step kill chains with MITRE ATT&CK mapping. Also handles general background tasks.

Reporting runs inline at scan completion — the scan worker assembles the final report at the end of the scan phase chain rather than dispatching to a separate worker. A dedicated reporting worker pool is planned for a future release as scan volume grows.

Data Stores

GHOST RED v2.5.0 uses SQLite for persistent storage in single-tenant deployments. This choice is deliberate: SQLite is fast, battle-tested, has zero operational overhead, and supports up to gigabyte-scale databases without issue. For multi-tenant and sovereign deployments, GHOST RED supports PostgreSQL as a drop-in replacement via SQLAlchemy ORM.

Intelligence Feeds

Two external intelligence feeds keep GHOST RED's knowledge of the threat landscape current:

- NVD (National Vulnerability Database) — synced continuously. Current production telemetry shows 1,250+ CVEs ingested in a representative 24-hour window.
- CISA KEV (Known Exploited Vulnerabilities) — 1,569 actively-exploited CVEs tracked as of April 2026, synced hourly. Findings that match KEV entries are flagged in scan reports as actively-exploited-in-the-wild, giving blue teams a prioritised view of what attackers are using right now.

Claude AI Integration

The AI Worker calls the Anthropic Claude API with a structured prompt containing: discovered CVEs, target topology, asset metadata, and scan phase context. Claude returns a structured JSON response representing a multi-step attack chain, with each step annotated by MITRE ATT&CK technique ID, likelihood score, and description.

PRODUCTION BENCHMARK

In a representative production run observed in v2.5.0 telemetry, a full AI chain generation — from discovered CVEs to a complete 3-step kill path mapped to MITRE ATT&CK — completed in 16.07 seconds. The resulting chain had a probability score of 0.75 and covered techniques T1595.001 (Active Scanning), T1190 (Exploit Public-Facing Application), and T1505.003 (Server Software Component / Web Shell).

Shipped Differentiators

Beyond the component list, GHOST RED v2.5.0 ships with four architectural properties that set it apart from both traditional scanners and first-generation AI security tools:

1. Isolation Between API and Execution

The Gunicorn application server and the Celery worker processes are fully decoupled. A restart, reload, or redeploy of the web layer does not kill scans that are in flight. Long-running scans continue on the worker pool until completion, and the UI reconnects to their state automatically on the next request. This is standard for mature SaaS but uncommon in early-stage security tooling.

2. Nested AI Reasoning Pattern

Scan tasks in the scan worker pool dispatch AI chain generation as separate tasks to the AI + Default worker pool. This means multiple scans can run in parallel without starving AI capacity, and a slow Claude API response on one scan does not block the next scan from starting. It also means the two pools can be scaled independently as customer load grows.

3. Single-Command Rollback Mode

The `SCAN_EXECUTION_MODE` environment variable lets operators switch between distributed Celery execution and a monolithic thread-based execution mode with zero downtime. If the Celery infrastructure has an issue, a single config change brings the platform back on a simpler architecture. This is a deliberate operational safety net for the single-founder stage.

4. Deep Health Monitoring

The `/api/health/deep` endpoint exposes the real state of the platform — not just a HTTP 200 heartbeat. It monitors six subsystems independently: database reachability, Redis connectivity, scan worker queue depth, AI worker availability, CISA KEV sync freshness, NVD sync freshness, and Claude API reachability. An UptimeRobot probe hits this endpoint every 5 minutes. Persistent cross-worker state is stored in a dedicated SQLite `app_meta` table so the health status survives worker restarts.

03 The 9 Scan Phases

How GHOST RED replicates the adversary lifecycle

GHOST RED models an attack the way a real adversary executes one — not as a list of checks, but as a sequential lifecycle where each phase informs the next. The 9 phases below mirror the MITRE ATT&CK enterprise matrix, adapted for autonomous execution.

Phase 1 — Reconnaissance

The platform maps the external attack surface: subdomains, exposed services, certificates, open ports, and public asset footprint. This phase is non-intrusive and uses a combination of active scanning, passive DNS enumeration, and certificate transparency log analysis.

MITRE mapping: T1595 (Active Scanning), T1590 (Gather Victim Network Information), T1589 (Gather Victim Identity Information)

Phase 2 — Enumeration

Services identified in Phase 1 are fingerprinted and version-detected. Web applications are crawled for directories, APIs, and authentication endpoints. Network services are probed for misconfigurations. This phase produces the asset inventory that drives the rest of the scan.

MITRE mapping: T1046 (Network Service Discovery), T1018 (Remote System Discovery), T1083 (File and Directory Discovery)

Phase 3 — Vulnerability Discovery

Every service identified is cross-referenced against the NVD and CISA KEV feeds. GHOST RED identifies both known CVEs (with CVSS scoring) and emerging threats (flagged by KEV presence). Custom checks supplement public data — for example, testing for weak TLS configurations, exposed secrets, and common misconfigurations.

MITRE mapping: T1595.002 (Vulnerability Scanning)

Phase 4 — Exploitation

For each exploitable finding, GHOST RED attempts safe validation — proof-of-concept execution that confirms the vulnerability without damaging the target. Destructive exploitation is gated behind explicit customer authorisation. This phase distinguishes GHOST RED from vulnerability scanners: exploitation validates which findings are real versus theoretical.

MITRE mapping: T1190 (Exploit Public-Facing Application), T1133 (External Remote Services), T1566 (Phishing — simulated only)

Phase 5 — Lateral Movement

Once initial access is established (or simulated), the platform maps potential lateral paths across the environment. This includes credential reuse analysis, trust relationship mapping, and internal network discovery. Lateral movement is simulated in audit mode by default — the platform reports possible paths without actually traversing them unless explicitly authorised.

MITRE mapping: T1021 (Remote Services), T1550 (Use Alternate Authentication Material)

Phase 6 — Privilege Escalation

GHOST RED identifies paths from low-privilege footholds to administrative access. This includes analysis of misconfigured service accounts, kernel exploits, sudo misconfigurations, and cloud IAM privilege escalation paths. The output is a ranked list of privilege escalation opportunities ordered by exploitability.

MITRE mapping: T1068 (Exploitation for Privilege Escalation), T1548 (Abuse Elevation Control Mechanism)

Phase 7 — Persistence

Simulated persistence mechanisms are identified and documented — places where a real attacker could maintain access. This includes account creation paths, scheduled task abuse, registry run keys, and startup folders on Windows environments. GHOST RED does not install persistence itself; it documents what a real attacker would do.

MITRE mapping: T1053 (Scheduled Task/Job), T1547 (Boot or Logon Autostart Execution), T1136 (Create Account)

Phase 8 — Exfiltration

The platform identifies what data an attacker could exfiltrate and by what means. This includes analysis of cloud storage permissions, database access paths, egress filter gaps, and covert channel opportunities. The output drives the data-classification component of the compliance report.

MITRE mapping: T1041 (Exfiltration Over C2 Channel), T1567 (Exfiltration Over Web Service)

Phase 9 — Impact

The final phase models what a successful attack would cost the business: encrypted data volume, services that would go offline, customer records exposed, regulatory fines triggered. This phase translates technical findings into business-language risk, which is what CISOs and boards need for decision-making.

MITRE mapping: T1486 (Data Encrypted for Impact), T1565 (Data Manipulation), T1498 (Network Denial of Service — modelled only)

> **SAFETY GUARDRAILS**
>
> Every phase that could cause damage — exploitation, lateral movement, persistence installation — is gated behind explicit customer authorisation levels. GHOST RED supports three execution modes: AUDIT (no active exploitation), VALIDATE (proof-of-concept only), and SIMULATE (full adversary emulation, for internal red team exercises in controlled environments).

Phase Execution Timing

A typical full scan — all 9 phases — completes in 45 to 90 seconds depending on target scope and authentication level. This is achievable because GHOST RED executes phases in parallel where it is safe to do so (for example, reconnaissance and enumeration can overlap across multiple targets) while preserving the sequential dependencies that matter (exploitation cannot start before vulnerability discovery has produced findings).

The 45–90 second figure is for the scan phases themselves. AI chain generation, which runs after scan completion in the AI + Default worker pool, adds approximately 14–16 seconds based on measured production telemetry.

04 AI Chain Engine

From isolated CVEs to executable kill paths

The fundamental limitation of traditional vulnerability scanners is that they produce lists. A list of 800 CVEs, each with a CVSS score, each rated independently. But real attacks do not work that way. Real attacks chain — a low-severity information disclosure feeds into a medium-severity authentication bypass which feeds into a critical remote code execution. The chain is what matters, not the individual CVEs.

GHOST RED's AI Chain Engine solves this problem. It takes the output of Phases 1 through 4 and asks the question: given what I have found, what is the most likely path a real attacker would take?

How Chain Generation Works

1. The scan worker finishes the scan phases and dispatches a chain-generation task to the AI + Default worker pool via Redis.
2. The AI worker constructs a sector-aware prompt for the Claude API. The prompt is specialised per target sector — Banking, Healthcare, Government, Enterprise — so the attack chain reflects realistic threat patterns for that industry (for example, financial fraud paths for banking, PHI exfiltration for healthcare).
3. The prompt includes: target environment summary, CVEs with CVSS scores and context, MITRE ATT&CK technique library reference, KEV correlation flags, and chain-generation guidelines.
4. Claude returns a structured JSON response containing one or more candidate chains, each with: a title, probability score (0.0 to 1.0), estimated execution time, and an ordered list of steps annotated with MITRE ATT&CK technique ID, CVE used (if applicable), target asset, technical description, and per-step likelihood score.
5. If the Claude API is unreachable (network failure, rate limit, outage), a rule-based fallback generator produces a chain from the discovered findings using heuristic rules — so every scan produces a chain, even in degraded API conditions.
6. The scan worker assembles the generated chain into the final report inline (reporting does not run in a separate worker in v2.5.0), producing both the technical report for the security team and the compliance evidence mapping for the auditor.

Sample Chain Output


```
{
  "chain_title": "Cloudflare Bypass to Web Shell",
  "probability": 0.75,
  "estimated_time": "2h 15m",
  "steps": [
    {
      "step": 1,
      "phase": "Reconnaissance",
      "technique": "Active Scanning",
      "technique_id": "T1595.001",
      "target": "wallettg.net:80,443,8080,8443",
      "description": "Enumerate subdomains, identify direct IP addresses
         bypassing Cloudflare protection, fingerprint backend
         technologies through header analysis and error pages",
      "likelihood": 0.9
    },
    {
      "step": 2,
```

```
      "phase": "Initial Access",
      "technique": "Exploit Public-Facing Application",
      "technique_id": "T1190",
      "target": "direct_ip:8080",
      "description": "Access unprotected backend server directly via
         discovered IP, exploit common web vulnerabilities (SQLi,
         file upload, deserialization) in wallet application",
      "likelihood": 0.8
    },
    {
      "step": 3,
      "phase": "Execution",
      "technique": "Server Software Component",
      "technique_id": "T1505.003",
      "target": "backend_server",
      "description": "Deploy encrypted web shell for persistent access",
      "likelihood": 0.7
    }
  ]
}
```


Why This Matters

A CVSS score tells you how severe a single vulnerability is. A chain tells you how close your environment is to a full compromise. This is the translation layer between technical findings and business risk — and it is what allows GHOST RED to produce compliance reports that map directly to DORA's operational resilience requirements or NIS2's incident-readiness mandate.

DEFENSIVE USE

Generated chains are not just offensive artefacts. Each chain includes recommended defensive controls — the specific detection rules, configuration hardening, and compensating controls that would break the chain. Blue teams use these to prioritise their defensive investments based on actual attacker behaviour, not theoretical CVE lists.

05 Compliance Framework Mapping

Audit-ready evidence, auto-generated

Every GHOST RED scan produces two report artefacts: the technical findings report (for the security team) and the compliance evidence package (for auditors and regulators). The compliance package is auto-generated by mapping every technical finding to the relevant framework control.

Live Frameworks (v2.5.0)

FRAMEWORK	SCOPE	COVERAGE
DORA	EU financial entities	Full TLPT mapping — Threat-Led Penetration Testing requirements under Article 26
NIS2	EU essential & important entities	Incident readiness evidence, risk management (Art. 21), supply chain (Art. 22)
ISO 27001	Information security management	A.5, A.8, A.12, A.13, A.14, A.16 controls with evidence artefacts
PCI DSS	Payment card industry	Req. 6 (secure development), Req. 11 (testing), Req. 12 (vulnerability management)

Roadmap Frameworks (2026)

FRAMEWORK	TARGET QUARTER	USE CASE
HIPAA	Q2 2026	US healthcare covered entities — administrative, physical, technical safeguards
NIST CSF 2.0	Q2 2026	US federal and commercial — Identify/Protect/Detect/Respond/Recover mapping
SOC 2	Q3 2026	SaaS and cloud service providers — Type II evidence collection
CMMC	Q3 2026	US defence industrial base — Levels 1 through 3 control mapping
SWIFT CSP	Q4 2026	Banking and financial messaging — Customer Security Programme controls

How Mapping Works

Each compliance framework is represented internally as a structured control library. Every technical finding produced by GHOST RED is tagged with the relevant control IDs at generation time. When the reporting worker compiles the compliance package, it performs a reverse lookup: for each control, what findings and remediations apply?

The output is a per-control evidence document that an auditor can consume directly. No manual translation needed. This is the step that traditionally costs consulting firms thousands of hours per year — and it is automated in GHOST RED.

AUDIT-READY OUTPUT

The compliance evidence package includes: control narrative, findings mapped to each control, remediation status, evidence screenshots, and a signed integrity hash that auditors can verify. Output formats: PDF (auditor-friendly), JSON (SIEM/GRC integration), and CSV (spreadsheet review).

Example: DORA Article 26 Mapping

DORA REQUIREMENT	GHOST RED EVIDENCE
Threat-Led Penetration Testing every 3 years	Continuous TLPT replaces point-in-time testing. Auto-generated test reports every scan.
Testing by qualified testers	GHOST RED operated by certified offensive security practitioners. TIBER-EU methodology alignment.
Coverage of critical functions	Asset criticality tagging. Crown jewel identification during Phase 1 reconnaissance.
Remediation follow-up	Automated re-testing of remediated findings. Closure verification in subsequent scans.
Executive and regulator reporting	Board-level risk summary + regulator-ready detailed evidence package in one output.

06 Deployment Models

SaaS, Sovereign, and Air-Gapped options

GHOST RED supports three deployment models, each designed for a different risk and compliance profile. The core platform code is identical across models; what changes is where data resides, how the platform is operated, and who has access.

Model 1 — SaaS (Standard)

The default deployment. Port Cyber Defense hosts the GHOST RED platform on EU-based infrastructure. Customers access via HTTPS, authenticate with SSO, and run scans against their environments. Scan results are stored in the customer's dedicated single-tenant instance.

CHARACTERISTIC	DETAIL
Host	EU-hosted infrastructure (region details available on request under NDA)
Data residency	EU-only by default, US region option planned as expansion milestone
Isolation model	Single-tenant per customer instance today. Multi-tenant architecture with per-tenant encryption keys is on the roadmap for H2 2026.
Pricing tier	Recon (€990/mo), Assess (€2,500/mo), Defend (€8,000/mo)
Ideal for	SMEs, mid-market enterprises, first-time buyers

Model 2 — Sovereign (On-Premise)

For organisations that cannot allow their security data to leave their own infrastructure, GHOST RED is available as a fully on-premise deployment. The customer hosts the platform on their own hardware or private cloud. Port Cyber Defense provides installation, configuration, updates, and support.

CHARACTERISTIC	DETAIL
Host	Customer-provided infrastructure (bare metal or private cloud)
Data residency	Entirely within customer's environment
Isolation model	Complete — no external dependencies except for CVE feed sync
Pricing tier	Sovereign License — from €150,000/year
Ideal for	Banks, government agencies, defence contractors, critical infrastructure

Model 3 — Air-Gapped

For the highest-security environments where no outbound connectivity is permitted, GHOST RED supports fully air-gapped operation. CVE feed updates are delivered via signed media (encrypted USB or one-way data diode). Scan results never leave the air-gapped network.

CHARACTERISTIC	DETAIL
Host	Customer-provided isolated network
Data residency	Zero external data flow
CVE feed updates	Delivered via signed encrypted media (monthly or weekly)
AI capability	Claude API access requires separate authorisation — offline AI mode

	available using local models
Ideal for	Military, intelligence, nuclear, critical defence contractors

DEPLOYMENT DECISION FRAMEWORK

If you are DORA-regulated but not a top-10 global bank: SaaS. If you are a bank with sovereign data requirements: Sovereign. If you handle classified data or critical national infrastructure: Air-Gapped. Most deployments sit in the SaaS tier; sovereign and air-gapped are designed for the 10% of customers with binding regulatory constraints.

07 Integration Roadmap

How GHOST RED fits into your existing stack

GHOST RED is designed to integrate with the security tooling enterprises already own. Every integration listed below is either live in v2.5.0 or scheduled for delivery in 2026. Our integration philosophy is outbound-first: GHOST RED produces structured outputs in standard formats (JSON, CEF, STIX/TAXII, PDF) that feed into existing SIEMs, ticketing systems, and GRC platforms.

SIEM Integrations

PLATFORM	STATUS	USE CASE
Splunk Enterprise	Q2 2026	Findings ingestion as CIM-compliant events. GHOST RED ships a certified app on Splunkbase.
Microsoft Sentinel	Q3 2026	Connector published on Sentinel Content Hub. Native KQL queries for chain analysis.
Elastic Security	Q4 2026	Ingest via Elastic Common Schema. Kibana dashboards bundled.
IBM QRadar	Q4 2026	DSM extension via QRadar SDK.

ITSM & Ticketing

PLATFORM	STATUS	USE CASE
ServiceNow	Q2 2026	Auto-create remediation tickets from findings. Custom table for CVE tracking.
Jira Service Management	Q2 2026	Direct webhook integration. Issue type templating by severity.
Freshservice	Q3 2026	API integration for mid-market customers.

Identity & Access

PLATFORM	STATUS	USE CASE
Azure Active Directory	Live (v2.5.0)	SAML 2.0 SSO. Group-based role mapping.
Okta	Live (v2.5.0)	SAML 2.0 SSO. SCIM provisioning available.
Google Workspace	Q2 2026	OIDC authentication.
Duo Security	Q3 2026	MFA enforcement for privileged actions.

Vulnerability Management

PLATFORM	STATUS	USE CASE
Tenable.io	Application in progress	Tenable Assure Partner application submitted. Integration design complete — bi-directional sync of vulnerability data scoped for delivery upon partner approval.
Qualys VMDR	Q3 2026	Import Qualys findings as GHOST RED scan

		inputs.
Rapid7 InsightVM	Q4 2026	API integration for joint customers.

Cloud & Marketplace

PLATFORM	STATUS	USE CASE
AWS Marketplace	Q2 2026	Self-service procurement with existing AWS budgets. Listed as SaaS subscription.
AWS Partner Network	Application submitted	APN application submitted March 2026. Benefits will activate upon partner approval.
Azure Marketplace	Q4 2026	Transact-enabled listing for Azure customers.
Google Cloud Marketplace	2027	Under evaluation.

08 2026 Product Roadmap

Quarter by quarter

The roadmap below is organised by quarter. Items marked LIVE are in production today. Items marked IN PROGRESS are actively being built. Items marked PLANNED are scheduled with clear ownership but not yet started.

Q1 2026 (Shipped)

- GHOST RED v2.5.0 shipped to production
- 9 scan phases operational
- 4 live compliance-mapping frameworks (DORA, NIS2, ISO 27001, PCI DSS)
- AI Chain Engine with MITRE ATT&CK mapping, sector-aware prompting, rule-based fallback
- Tenable Assure Partner application — submitted
- AWS Partner Network application — submitted
- Enterprise whitepaper published

Q2 2026 (In Progress)

- HIPAA compliance framework — target delivery June
- NIST CSF 2.0 mapping — target delivery June
- Splunk Enterprise certified app (Splunkbase)
- ServiceNow and Jira ticketing integrations
- AWS Marketplace SaaS listing
- First SaaS client pilot (target: 1–2 enterprise customers)
- Seed round close (€300k target)

Q3 2026 (Planned)

- SOC 2 compliance framework
- CMMC compliance framework
- Microsoft Sentinel connector
- Qualys VMDR bi-directional integration
- Offline AI mode for air-gapped deployments
- Hack The Box partnership activation (talent + training)
- 3–5 paying SaaS clients target

Q4 2026 (Planned)

- SWIFT CSP compliance framework
- Elastic Security + IBM QRadar connectors
- Azure Marketplace listing
- Rapid7 InsightVM integration
- Multi-region deployment (US availability)
- Series A preparation and pitch

Stretch & 2027 Items

- Real-time continuous scanning (shift from scheduled to always-on)
- Purple team mode — coordinated defensive exercises with blue team tooling
- Custom compliance framework builder for regional/industry standards
- LATAM market expansion (Mexico, Colombia, Chile)
- Google Cloud Marketplace

ROADMAP CAVEAT

Roadmap items and dates reflect current planning as of April 2026. Seed funding outcomes, early customer feedback, and partnership progress will shape prioritisation. Commitments become binding only on explicit written agreement with Port Cyber Defense.

09 Security & Data Handling

How we protect what you entrust to us

GHOST RED is a security product. We hold ourselves to a higher bar than general-purpose SaaS because the data we handle — vulnerability findings, attack chains, topology maps — would be extraordinarily valuable to a real adversary. Our security posture reflects that.

Data Classification

DATA TYPE	HANDLING
Customer scan results	Stored in the customer's single-tenant SQLite database instance. Encryption at rest is provided by the host filesystem's disk-level encryption. Retention: 12 months default, configurable down to 30 days.
Customer credentials (for authenticated scans)	Stored in a restricted-permission secrets file on the host (0640, root:service-user). Accessible only to the scan worker process at execution time.
AI prompt contents	Minimised — only sanitised findings metadata is sent to the Claude API. No customer credentials, no raw PII, no full network captures.
Attack chain output	Treated as sensitive. Delivered to customer only via authenticated portal or encrypted email.
Customer account metadata	Standard SaaS practice. GDPR-compliant handling. Right to deletion honoured within 30 days.

CURRENT STAGE HONESTY

GHOST RED v2.5.0 ships as a single-tenant deployment — each customer instance is isolated from every other. Multi-tenant architecture with per-tenant encryption keys is on the roadmap for H2 2026 and will be delivered before the first enterprise tier customer that requires it. We describe what we ship, not what we plan.

Access Controls

- Role-based access control (RBAC) with four default roles: Viewer, Operator, Admin, Owner
- Mandatory MFA for Admin and Owner roles
- SAML 2.0 SSO integration (Azure AD, Okta)
- All privileged actions audit-logged with user, timestamp, and source IP
- Session timeout: 15 minutes idle, 8 hours absolute maximum
- Founder-level root access to production is the standard exception during the current product stage and is audit-logged. Separation of duties with a dedicated operations role is planned following seed close.

Infrastructure Security

- Network isolation: private subnets, no direct internet access for scan workers
- Egress filtering: only whitelisted endpoints (NVD, KEV, Claude API, customer-authorised scan targets)
- Secrets management: filesystem-level access control today (restricted permissions, service-user ownership). Migration to HashiCorp Vault or AWS Secrets Manager is planned post-Series A.
- Monitoring: all services emit structured logs. The /api/health/deep endpoint reports six-subsystem status, polled every 5 minutes by UptimeRobot.
- Incident response: documented playbooks and founder direct-contact. Business-hours support (CET) with next-business-day acknowledgement for standard issues, same-day acknowledgement for critical issues. 24/7 on-call rotation with 4-hour SLA is on the roadmap following seed close and first-hire onboarding.

Compliance Mapping vs. Certification

A deliberate note on language. GHOST RED maps scan findings to the control requirements of DORA, NIS2, ISO 27001, and PCI DSS — the platform produces evidence packages structured around these frameworks. This is compliance mapping.

Port Cyber Defense itself is not currently certified to any of these standards. ISO 27001 certification and SOC 2 Type I are on the 2027 roadmap following organisational maturity required to achieve them. We will not describe the company as certified until certificates are issued.

Certification Roadmap

- GDPR compliant — full data subject rights supported today
- ISO 27001 certification — targeted 2027 at earliest, dependent on team size and documented management systems
- SOC 2 Type I — targeted 2027, with Type II following once the required 6–12 month observation window is complete

RESPONSIBLE DISCLOSURE

If you identify a security issue in GHOST RED itself, we operate a responsible disclosure programme. Contact info@port-cyber-defense.com with details. We acknowledge within 24 hours and commit to a fix timeline based on severity. Researchers are credited publicly (with permission).

10 Team & Contact

Who builds this, and how to reach us

Founder

AHMAD GHALEB

Founder & CEO

Ahmad founded Port Cyber Defense in 2023 after a decade of offensive security practice. He built the GHOST RED platform from first principles with the conviction that autonomous AI adversary simulation would replace manual red teaming within five years. Ahmad has personally delivered penetration testing engagements for enterprise clients across banking, finance, and healthcare before pivoting full-time to platform development in 2025.

Hiring Plan

Following the seed round close (Q2 2026 target), Port Cyber Defense will hire two initial engineers, both sourced from the Hack The Box ranked community:

- Senior Offensive Security Engineer — to expand the scan worker capabilities and lead exploit development
- AI/ML Engineer — to evolve the Chain Engine beyond Claude API into bespoke fine-tuned models

Both hires are targeted for Q3 2026. Long-term team plan: 5–10 people by end of 2026, 15–25 by end of 2027.

Contact

GENERAL INQUIRIES	info@port-cyber-defense.com
Technical questions	Submit via /technical-brief form — response within 24 hours
Partnership discussions	info@port-cyber-defense.com
Security issues	info@port-cyber-defense.com (responsible disclosure)
Press & media	info@port-cyber-defense.com
Website	port-cyber-defense.com
LinkedIn	linkedin.com/company/port-cyber-defense
Offices	Málaga, Spain • Strasbourg, France

END OF TECHNICAL BRIEF