CISA KEV · FIRST.org EPSS · OpenTelemetry / OTLP · Splunk HEC

Security operations built for the way your infrastructure actually breaks.

Obiguard SOC streams your logs, metrics, traces, dependencies, and commits into one place — and turns them into alerts with AI-written executive summaries, CVEs cross-matched against real-world exploit intel, and predictive risk scores before the error rate itself trips an alarm.

- *Detection runs out of the box — no rules engine to configure before day one.*

10	2	24/7	0
MONITORING & CONFIG SURFACES	THREAT FEEDS — CISA KEV + FIRST EPSS	CONTINUOUS DETECTION	DETECTION RULES TO WRITE

The platform, not the pitch deck.

9

Monitoring Surfaces

Alerts, Errors, CVE Radar, Threat Intel, Precog, Live Logs, Metrics, Traces & Service Map — one pane of glass.

2

Threat Intelligence Feeds

Every CVE finding is cross-matched live against CISA's KEV catalog and FIRST.org's EPSS exploit-prediction score.

0

Dashboards To Build

Metrics, traces, and the service map render themselves the moment data starts flowing in — nothing to configure.

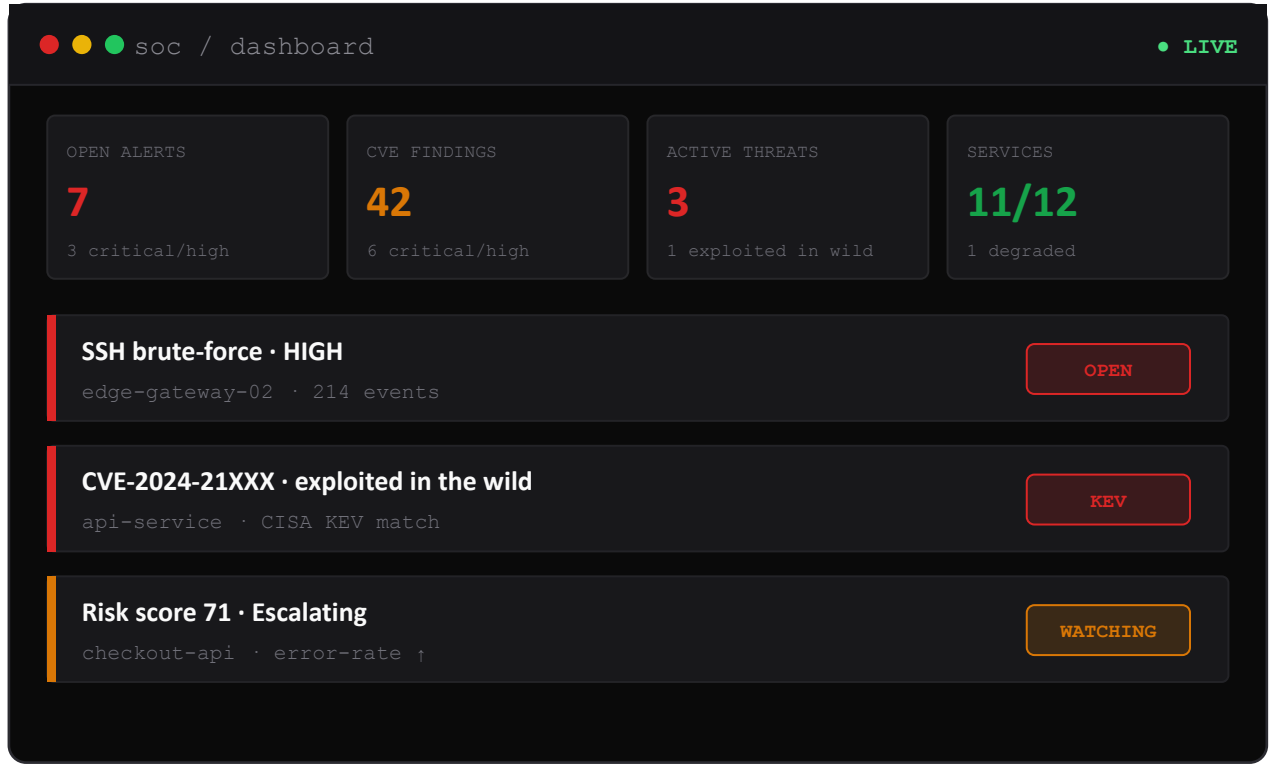

Actual SOC Dashboard view: alerts, CVE findings, threat intel, and fleet health rendered on one live page.

A live view from a customer's production environment.

3 alerts fired in the past hour.

1 flagged by Threat Intelligence as actively exploited — CISA KEV.

2 CVE findings auto-triaged with an upgrade safety score.

Every one of them is written to the evidence timeline — searchable, exportable, and ready for the next audit.

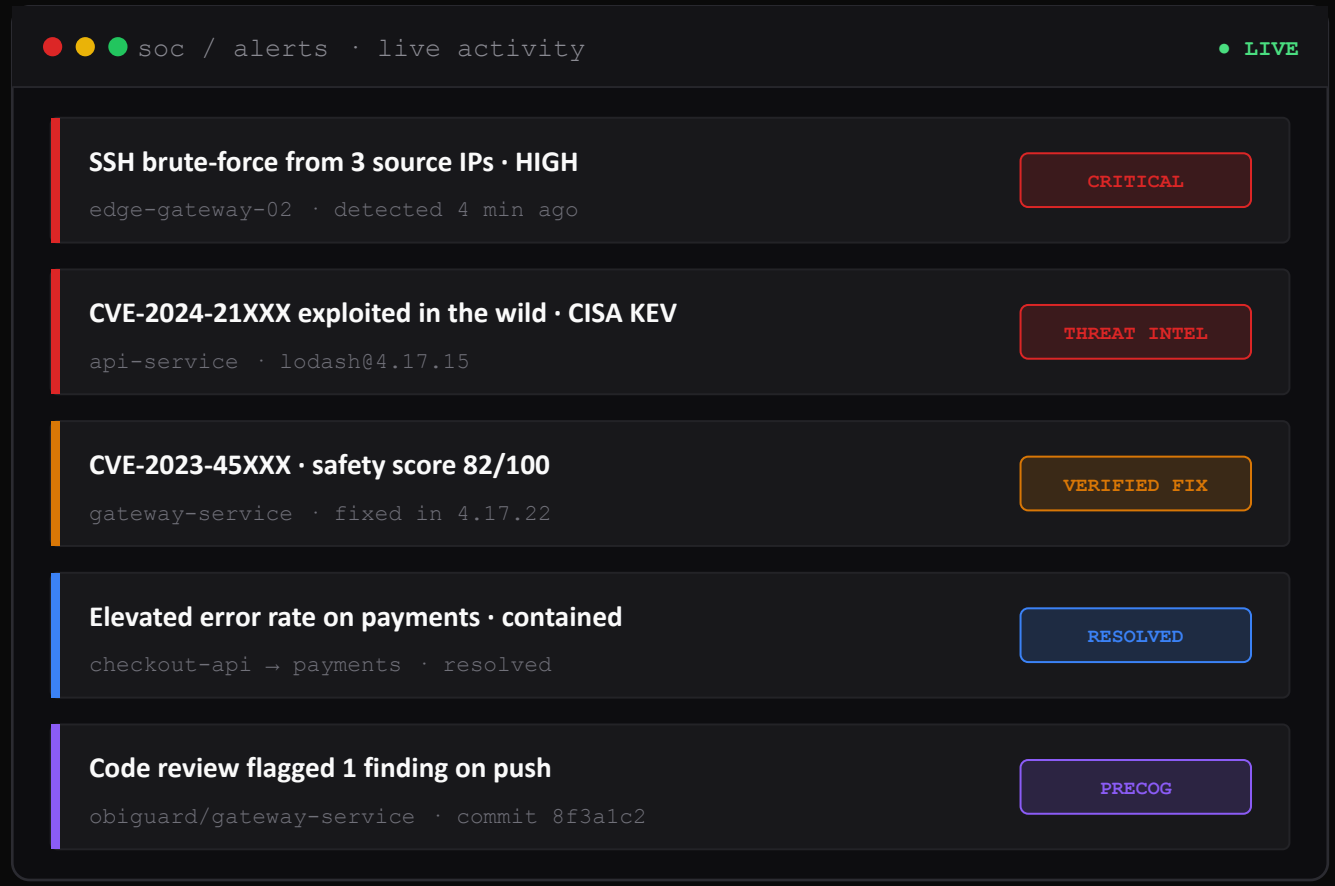

Six capabilities. One security operations center.

Obiguard SOC wraps every log, metric, trace, dependency, and commit your organization produces — detecting, explaining, and tracking issues without a rules engine to configure.

01

MONITOR

Alerts & Incident Response

Detected automatically from log volume and rule matches. Each incident gets an AI executive summary and a tiered remediation plan.

02

TRACK

Errors

Application errors auto-grouped by fingerprint across service and environment, with symbolicated stack traces.

03

SCAN

CVE Radar & Threat Intel

Every push scanned for vulnerable dependencies, then cross-matched against CISA KEV and FIRST.org EPSS.

04

PREDICT

Precog

Predictive risk scoring from log and error-rate trends, plus an automatic LLM code review on every commit.

05

OBSERVE

Logs, Metrics, Traces & Service Map

Streaming logs, host metrics, and distributed traces — with a service map that builds itself from what's running.

06

MEASURE

Code Coverage

Track per-repo test coverage from any lcov-emitting tool, uploaded with a one-line GitHub Action.

Built for the whole room — not just the SOC.

Security, platform, and engineering teams all look at the same data today — through different tools. SOC gives each of them the surface they actually need.

FOR SECURITY

Investigate an incident without hopping between five tools.

- Live alert queue with AI executive summaries & business-impact assessments
- Full evidence timeline linked to the raw log events that triggered it
- Errors auto-grouped by fingerprint — one bug shows up once, not a thousand times

For SOC analysts & incident responders

FOR PLATFORM & SRE

See what's actually running, and what's currently broken.

- Live Logs, Metrics, and Traces in one place — no dashboards to build
- Service map builds itself automatically from ingested traces
- Silent-host and degraded-service detection, per host and per edge

For platform, infra & SRE teams

FOR APPSEC & ENGINEERING

Fix what matters first, and ship with confidence.

- CVE Radar scans every push + daily, with an upgrade safety score
- Threat Intel flags which open CVE is being exploited in the wild right now
- Coverage trend and an automatic LLM code review on every commit

For AppSec & engineering teams

How it works: ingestion architecture.

Logs, metrics, traces, and coverage all land through the same Obiguard gateway your LLM traffic already uses — detection and storage run behind it.

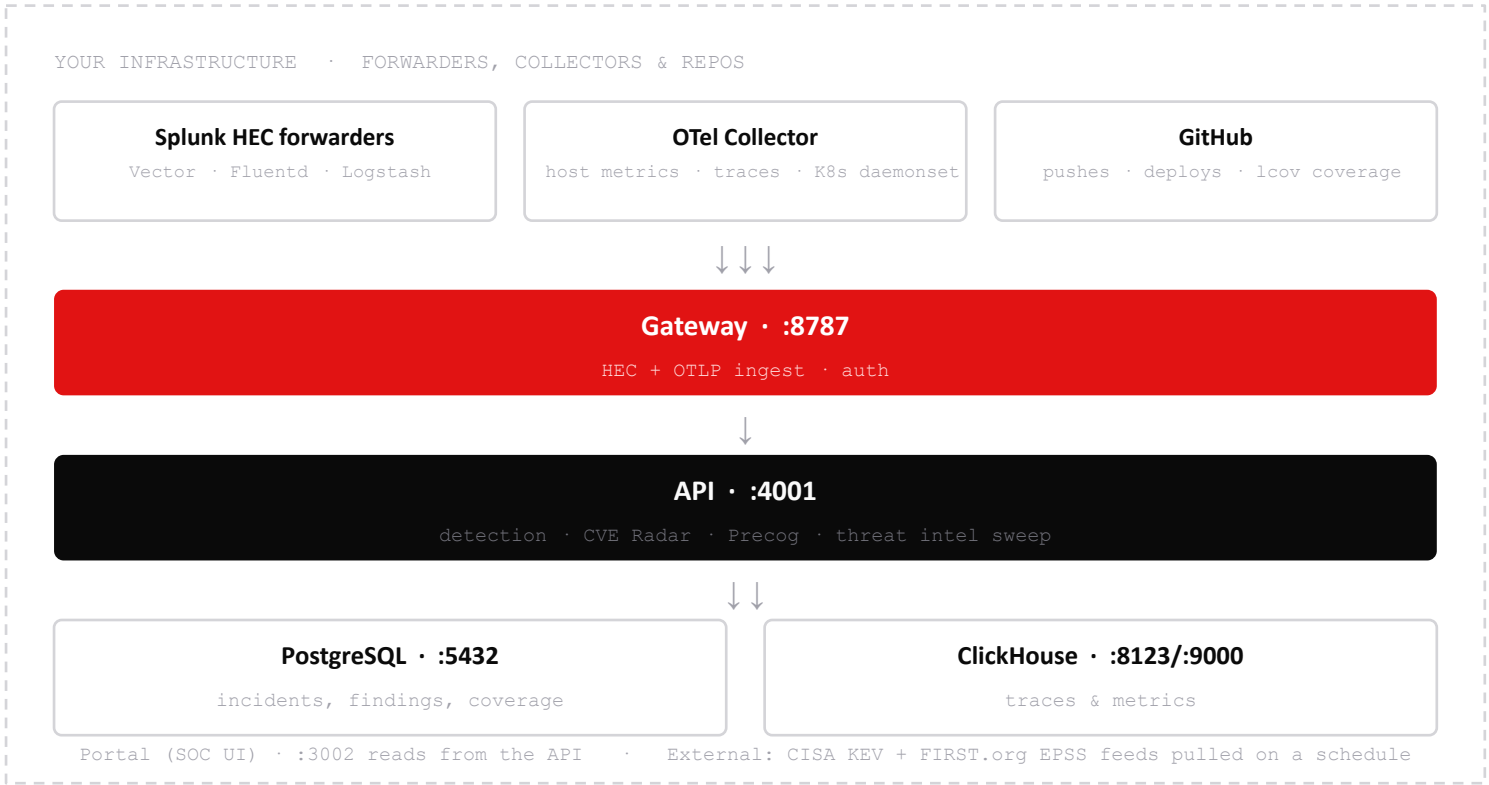

Same gateway, one ingest path

Logs, metrics, traces, and coverage all land through the same Obiguard gateway your LLM traffic already uses — no separate collector stack to run.

Scheduled + on-push detection

CVE Radar scans on every push and once a day. Threat Intelligence sweeps CISA KEV and FIRST.org EPSS on a schedule — no cron jobs to babysit.

Kubernetes in one command

A ready-made OTel Collector daemonset covers EKS, GKE, and AKS — host metrics and traces start flowing without hand-built config.

Common questions.

Don't see what you're looking for? Our solutions engineers respond within one business day.

01

Which connectors are actually live today?

GitHub is fully connected — deployments, pushes, workflow runs, and repo scanning. Kubernetes (EKS/GKE/AKS) ships as a one-command OTel Collector daemonset. Log ingestion works with any Splunk HEC-compatible forwarder, and coverage uploads work from any CI via our GitHub Action.

02

How is Threat Intelligence different from CVE Radar?

CVE Radar finds vulnerable dependencies in your repos and tells you how to fix them safely. Threat Intelligence cross-references those findings against CISA's KEV catalog and FIRST.org's EPSS score, so you know which one is being exploited in the wild right now.

03

Do I need a separate APM tool for traces and service maps?

No. Traces ingest over standard OTLP through the same collector you run for host metrics, and the Service Map builds itself automatically from whatever traces come in — no topology to draw by hand.

04

Do I need to write detection rules myself?

No. Alerts fire from log-volume and rule matches out of the box, errors auto-group by fingerprint, CVE Radar scans on every push and daily, and Precog scores risk continuously — you connect sources and get findings.

See what's actually happening in your infrastructure — before your customers do.

A 20-minute call is enough to scope your first connector. Most teams are streaming logs and metrics within a day, and enforcing detections in production within a week.

Email: hello@obiguard.com →

https://obiguard.ai/soc ↗

Security operations for the infrastructure your AI — and everything else — actually runs on.