OFFLOAD SECURITY

Threat Intelligence

Module Whitepaper

Real feeds, IOC correlation, and MITRE ATT&CK context on your own findings

Version 1.0 · July 2026 · Classification: Customer-Shareable

1. Overview

The Threat Intelligence module ingests indicators from nine external providers, manages IOCs with confidence and aging, models actors and campaigns in STIX 2.1, and — most importantly — correlates threat context onto the platform's own findings to sharpen prioritization.

2. Key Capabilities

- Nine live feed providers: CISA KEV, AlienVault OTX, Abuse.ch, Feodo Tracker, SSL Blacklist, PhishTank, Blocklist.de, Spamhaus DROP, and OpenPhish — real ingestion and parsing, per-team feed management.
- Per-team encrypted API keys for authenticated feeds.
- IOC management: indicators with confidence, severity, and aging; IOC-to-finding correlation.
- Vulnerability prioritization with threat context: known-exploited and actively-abused signals applied to the platform's own vulnerability queue.
- Threat actors and campaigns modeled in STIX 2.1 with kill-chain, TLP, and MITRE ATT&CK mapping.
- MITRE ATT&CK heatmap built from ingested actor techniques.
- Alerting rules, threat-hunting queries, and landscape reports.

3. How It Works

Feeds are pulled through a provider factory with real HTTP ingestion and per-provider parsers; failures are classified (auth, rate-limit, format) rather than silently swallowed. Indicators are stored team-scoped with confidence scoring and aging so stale intelligence decays.

Correlation runs against the platform's unified findings: an IOC or KEV match on one of your vulnerabilities raises its priority in the same queue your team already works from — threat intelligence as an input to operations, not a separate console.

4. Data Handling & Security

Aspect	How it is handled
Credential protection	All connection credentials are encrypted at rest with Fernet authenticated encryption (AES-128-CBC + HMAC), decrypted in memory only at the moment of use, and never returned by APIs.
Tenant isolation	Every record is scoped to the owning team; all reads and writes filter on the tenant identifier, enforced by database indexes.
Retention & deletion	Results are subject to configurable retention windows; data can be hard-deleted per scan and is removed by the tenant-offboarding cascade.
Deployment	Available as SaaS or fully self-hosted / on-premises — in the on-premises model no

Aspect	How it is handled
	security data leaves your environment.

5. One Platform, One Risk View

Threat context enriches unified vulnerability management and the risk register, powers the MITRE heatmap, and can trigger central alerts; feed health is monitored like any other platform job.