Threat Teaming Framework: Operationalizing Offensive Intelligence

Closing the Gap Between Detection and Remediation via Automated Threat Modeling

Publish Date: July 2026 **Classification:** Public

https://threatteamingframework.com
https://threatteamingframework.com/demo

A Hoffmann Holdings Asset
hello@hoffmann.holdings
https://hoffmann.holdings

Executive Summary

The "Mean Time to Remediate" (MTTR) for critical exposures is the defining metric of defensive failure. Organizations spend millions on offensive assessments, yet the eventual intelligence and PDF reports remain unstructured, resulting in effectively dead data. This latency creates a window of exposure that adversaries exploit with increasing velocity.

The Threat Teaming Framework (TTF) is a Continuous Threat Exposure Management (CTEM) platform that automates the ingestion, analysis, and operationalization of offensive security data. It replaces manual spreadsheet tracking with a deterministic, API-driven pipeline, and extends beyond exposure management into incident response with automated regulatory deadline tracking, detection engineering, vulnerability lifecycle management, and a full executive (virtual-CISO) toolkit, all synthesized into a single Command Center.

Strategic Imperatives:

- **Kill "Scan and Pray":** Shift from volume-based vulnerability management to threat-informed prioritization based on actual adversary behavior.
- **Quantify Negligence:** Replace subjective "High/Medium" ratings with a Rhino Score (0-850) underpinned by Monte Carlo financial simulations, directly translating technical debt into expected financial loss.
- **Operationalize Intelligence:** Aggregate 10 threat feeds into a unified operational view, reducing external API costs by >90% via intelligent caching while providing real-time context on IOCs.
- **Automate Compliance Evidence:** Deterministically map technical findings to PCI-DSS v4.0, NIST 800-53, and CIS v8 controls, transforming every remediation ticket into audit evidence.
- **Active Defense:** Move beyond passive reporting to active Purple Team orchestration, validating detection rules against known APT techniques.
- **Respond with the Clock Running:** When an incident is declared, TTF automatically determines which notification regulations apply (GDPR, SEC, HIPAA, state breach laws) and counts down every deadline with a live war room, AI-drafted notifications, and post-mortems.
- **Govern from One Pane:** The CISO Command Center synthesizes posture, ranked next actions, and 30-day trends across every module into a single score, with board-ready PDF summaries on demand.

Table of Contents

---

1. The Problem: The Offensive Intelligence Gap

Organizations are currently trapped in a cycle of "Intelligence Decay." High-value findings from elite Red Team engagements are effectively buried in a "PDF Black Hole." Unstructured, non-searchable reports represent millions in wasted security investment. This Offensive Intelligence Gap creates a window of liability where adversaries exploit known vulnerabilities while defenders manually copy data into spreadsheets. TTF immediately recovers this wasted ROI by converting dormant PDF data into a live, API-driven intelligence asset.

- **Unstructured Data Liability:** High-value findings from Red Team engagements are trapped in PDFs. Manual extraction is slow, lossy, and unscalable.
- **Context Switching Costs:** Analysts waste ~40% of their time correlating a finding (Jira) with threat intel (VirusTotal/Shodan) and compliance impact (GRC tools).
- **Remediation Regression:** Without a unified history of "fixed" techniques, organizations repeatedly pay pentesters to find the same vulnerabilities.
- **Audit Paralysis:** Mapping technical flaws to governance frameworks is a manual, quarterly panic rather than a continuous, automated process.
- **CISO Criticality:** Every day that a critical finding sits in a PDF inbox is a day of accepted liability. TTF eliminates this latency.

2. Threat Model: Adversary Economics

Adversaries operate on a graph; defenders operate on a list. TTF aligns defensive operations with adversary reality by modeling:

- **Identity as Perimeter:** Attackers abuse valid credentials for lateral movement. TTF tracks compromised identities via HIBP integration.
- **TTP Reusability:** APTs (e.g., APT29, Lazarus) reuse tools and techniques. If you cannot detect a T-Code (e.g., T1059) seen in a previous report, you are vulnerable to the entire campaign.
- **Living off the Land:** Attacks using native binaries (LOLBAS) evade signature detection. TTF's AI parser identifies these nuanced behaviors from report narratives.
- **Ransomware as a Service (RaaS):** The ransomware economy has industrialized. Affiliates purchase access, deploy payloads, and split profits with operators. TTF maps your exposure against 74+ tracked ransomware groups, 55 techniques across 11 kill chain phases, and profiles of the 20 most active operators including LockBit, BlackCat, Cl0p, Akira, and Black Basta.
- **Vulnerable Driver Exploitation:** Attackers increasingly abuse signed kernel drivers to disable EDR and gain ring-0 access. TTF cross-references every report against the LOLDrivers catalog of vulnerable and malicious drivers, flagging findings that indicate kernel-level compromise paths.

3. The Platform: Threat Teaming Framework

TTF is the middleware between offensive identification and defensive remediation. It creates a single source of truth for exposure management. TTF acts as a SaaS platform that ingests unstructured offensive reports, enriches them against 10 curated intelligence sources and 290+ threat actor profiles, maps findings to 11+ GRC frameworks, scores financial risk via Monte Carlo simulation, and generates actionable remediation plans - all within minutes of upload.

Beyond individual report analysis, TTF provides ransomware readiness intelligence, 13 purpose-built novel analytics engines, a CISO Command Center, full incident response with automated regulatory deadline tracking, a detection engineering advisor with SIEM alert triage, vulnerability lifecycle management with SLA enforcement, an executive toolkit (budget memos, board meeting prep, peer benchmarking, insurance gap analysis, M&A due diligence), campaign tracking, critical findings alerts, scheduled reporting, and a citation-backed natural language AI assistant. It is designed to operationalize the five stages of the Gartner CTEM cycle:

1. **Scoping:** Defining attack surfaces of what was assessed during your Red Team or Penetration Test.
2. **Discovery**: Automated ingestion and entity extraction from all offensive assessment history.
3. **Prioritization**: Threat-informed risk scoring using Monte Carlo financial modeling rather than simple severity ratings.
4. **Validation**: Active verification of defensive controls via the Purple Team Exercise Platform and Atomic Red Team tests.
5. **Mobilization**: Direct integration with Jira, GitHub, and SIEM/SOAR platforms to drive remediation velocity.

4. Architecture & Data Flow (Deep Dive)

TTF employs a serverless, event-driven architecture optimized for data isolation and asynchronous processing.

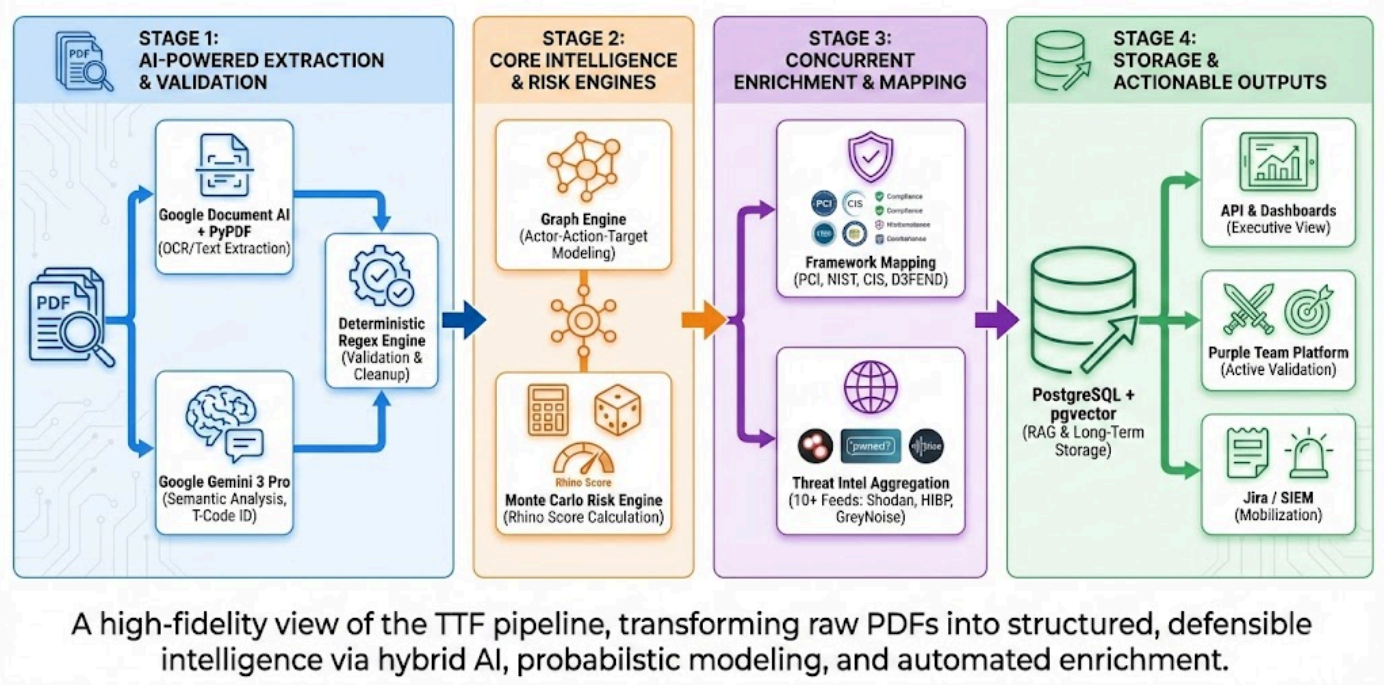

Figure 1: The ingestion pipeline

Reports are ingested via a secure upload stream to Google Cloud Run. The AI Extraction Layer (Python/FastAPI) utilizes pdfplumber for text extraction and Gemini 3 Pro (1M+ context) for semantic analysis, with a regex safety net. Data is normalized and passed to the Enrichment Engine, which queries the PostgreSQL Threat Intel Cache (Shodan/HIBP/ThreatFox). The fully hydrated object is stored in Cloud SQL (PostgreSQL 14+) using pgvector for semantic search. The Risk Engine triggers a Monte Carlo simulation (10k iterations) to update the Rhino Score. Finally, the API Layer exposes actionable data to SIEM/SOAR integrations.

Technical Stack Specifications

- **Runtime:** Python 3.11 (FastAPI) on Google Cloud Run.
- **Database:** PostgreSQL 14+ with pgvector and JSONB for flexible schema handling.
- **AI/ML:** Google Gemini via the google.genai SDK, paired with deterministic validation layers.
- **Asynchronous Task Queue:** Cloud Tasks for long-running ingestion and validation jobs.
- **Edge Protection:** Global external load balancer with Cloud Armor (per-IP rate-limiting bans and WAF rulesets) in front of all API traffic; Cloud Run ingress locked to the load balancer.
- **Security Middleware:** CSRF (Double Submit Cookie), Rate Limiting (Token Bucket), Security Headers (HSTS/CSP).

5. Core Mechanisms (How It Works)

1. Deterministic T-Code Extraction (Hybrid AI/Regex)

Unlike generic LLM wrappers, TTF uses a hybrid approach. Gemini 3 Pro extracts semantic context (reasoning, evidence), while a strict regex engine (_TID_RE) validates MITRE ATT&CK Technique IDs (Txxxx) [1]. This prevents hallucinated techniques and ensures data integrity for downstream mapping.

2. Monte Carlo Risk Simulation

The Rhino Score is not an arbitrary number. It is derived from a Monte Carlo simulation (10,000 iterations) that calculates "Expected Loss" [2]. The model inputs include a baseline Incident Response cost ($56,000), finding confidence, and user calibration scores. It applies probability multipliers based on disposition (Executed vs. Observed) to model financial exposure variance.

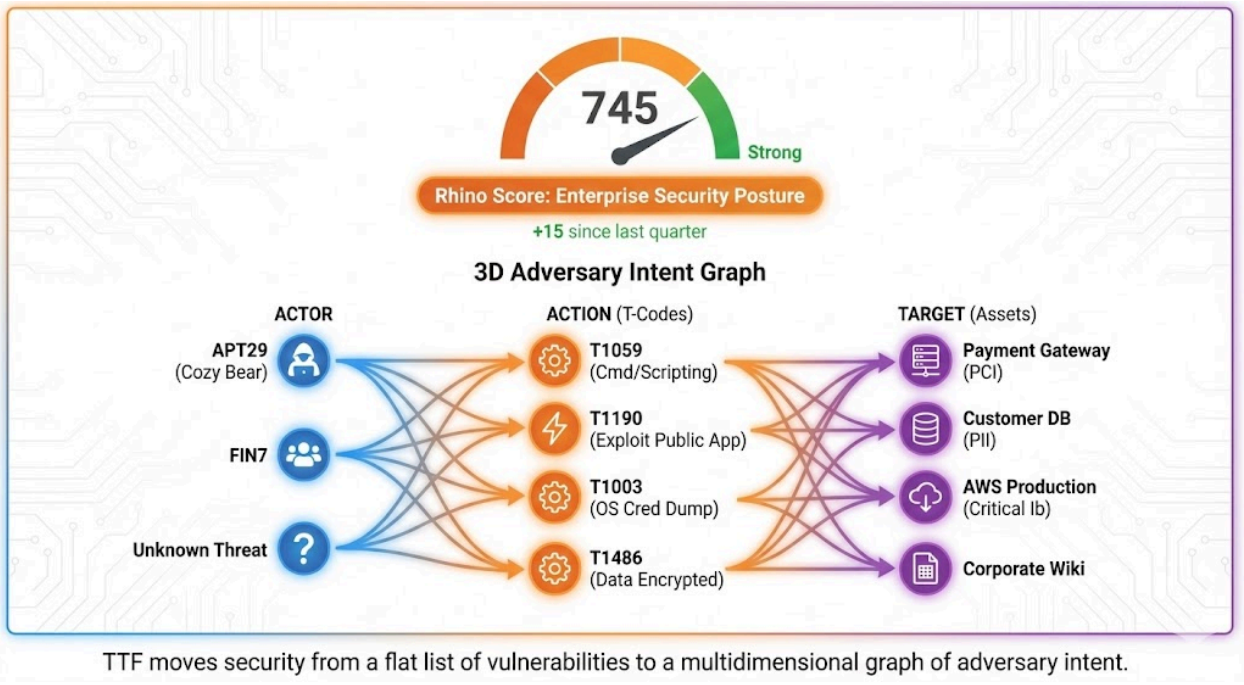

Figure 2: Rhino Score based on Monte Carlo simulations using multiple variables

Executive Spotlight: The Common Language of Risk

The Rhino Score (0-850) provides security leadership with what has been missing for a decade: a common language for the Board of Directors. Much like a FICO credit score translates complex financial history into a single, actionable metric, the Rhino Score distills 10,000+ Monte Carlo simulations and technical debt into a quantifiable measure of defensive health. It shifts the conversation from abstract "High/Medium/Low" findings to a defensible, data-driven posture assessment that supports strategic investment and insurance qualification.

3. Automated Cross-Walking (Enrichment Engine)

The enrichment service automatically maps ingested findings to compliance frameworks such as PCI-DSS v4.0 and NIST 800-53 [3]. A finding mapped to a MITRE Technique is instantly cross-referenced against CIS Controls v8, NIST 800-53, PCI-DSS v4.0, and D3FEND countermeasures.

This logic is hard-coded in the enrichment modules ensuring consistent, audit-ready mappings.

4. Smart Threat Intelligence Caching

To provide enterprise-grade intelligence without enterprise-grade costs, TTF implements a "Smart Caching" layer for IOC and breach data lookups [4]. IOC lookups (IP/Hash) are cached for 24 hours; breach data (Email/Domain) is cached for 7 days. This creates a local threat intelligence database that reduces external API calls (and costs) by over 90% while maintaining operational relevance.

5. Vector-Based Contextual Search

The database utilizes pgvector to store embeddings of report content (`DocumentChunk`). This allows analysts to perform semantic searches across historical reports (e.g., "Show me all past instances of SQL injection in payment gateways") rather than simple keyword matches, identifying systemic patterns.

6. Automated Threat Modeling (STRIDE / DREAD / PASTA)

Every uploaded report is automatically assessed against three industry-standard threat modeling frameworks. STRIDE classifies threats by category (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege). DREAD scores each technique on Damage, Reproducibility, Exploitability, Affected Users, and Discoverability. The PASTA (Process for Attack Simulation and Threat Analysis) 7-stage model provides business-context risk assessment. All three models are computed deterministically - no manual analyst input required.

7. Critical Findings Alerts (Compound Evidence)

TTF includes a compound-evidence alerting system designed to eliminate false positives and alert fatigue. Rather than firing on every "Critical" finding, the system evaluates six independent signal types:

- **CISA KEV Match:** Finding exploits a known-exploited vulnerability.
- **Detection Blind Spot:** No Sigma or Elastic rule covers the technique [5].
- **Critical DREAD Score:** DREAD total exceeds the critical threshold.
- **Ransomware Kill Chain:** Technique maps to an active ransomware group's playbook.
- **Vulnerable Driver:** Finding involves a known-malicious kernel driver (LOLDrivers).
- **APT Attribution:** Technique is associated with a tracked threat actor group.

An alert fires only when two or more independent signal types converge on the same finding with a combined weight exceeding the threshold. This ensures that only genuinely impactful findings trigger email notifications to the report uploader. Each alert includes a specific, actionable recommendation tailored to the combination of signals detected.

6. Ransomware Readiness Intelligence

Ransomware is the single largest financial threat facing most organizations. TTF provides a purpose-built ransomware intelligence engine that goes far beyond simple vulnerability scanning.

Coverage and Scope

TTF tracks 55 MITRE ATT&CK techniques across 11 kill chain phases specific to ransomware operations: Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, Collection, Exfiltration, and Impact. Each phase is weighted by operational significance.

The platform maintains enriched profiles of 20 high-activity ransomware operators (LockBit, BlackCat/ALPHV, Conti, Black Basta, REvil, Cl0p, Play, Royal, Akira, Medusa, Rhysida, BianLian, Hunters International, Scattered Spider, 8Base, Vice Society, Hive, Cuba, BlackByte, and INC Ransom). An additional 54+ groups are tracked via the GapMATRIX integration, bringing the total to 74+ monitored ransomware groups.

What TTF Computes

For every analyzed report, the ransomware engine calculates:

- **Kill Chain Phase Coverage:** How many of the 11 phases are addressed by your detected techniques.
- **Group Exposure Profile:** Which specific ransomware groups your environment is most vulnerable to, based on TTP overlap.
- **Initial Access Broker (IAB) Risk:** Whether detected techniques indicate the involvement of access brokers who sell footholds to ransomware affiliates.
- **Double/Triple Extortion Risk:** Analysis of data staging, exfiltration pathways, and evidence of extortion escalation patterns.
- **Financial Impact Estimate:** Projected ransom demand, operational downtime cost, recovery cost, and regulatory penalty exposure - calibrated by industry vertical.
- **Recovery Complexity Score:** Assessment of backup resilience, encryption speed profiling, and estimated recovery timeline.
- **Purple Team Exercises:** 8 pre-built exercise templates mapping directly to the most critical ransomware techniques, each with tool recommendations and detection validation steps.

The CISO receives a single Ransomware Readiness Score (0-100) with a breakdown by kill chain phase, matched group profiles, and prioritized remediation actions.

7. Threat Actor Intelligence (GapMATRIX)

Understanding who is targeting you is as important as understanding what they use. TTF integrates with the GapMATRIX community intelligence database to provide real-time threat actor correlation [6].

290+ Threat Actor Profiles

Every analyzed report is automatically matched against 290+ threat actor profiles encompassing nation-state APTs, financially motivated groups, hacktivists, and ransomware operators. Matching is based on TTP overlap - not simple keyword searches - ensuring high-confidence attribution.

300+ CVE-to-Actor Mappings

TTF maintains a curated database of which threat actors exploit which CVEs. When a report identifies a vulnerability, the platform immediately answers: "Which threat groups are known to exploit this CVE?" and "What other techniques do those groups use?" This enables defenders to anticipate the next steps of an adversary already inside their environment.

Automated Weekly Sync

Threat actor data is automatically refreshed weekly via a dedicated CI/CD pipeline. New actor profiles, TTP updates, and CVE correlations are ingested without manual intervention, ensuring the intelligence layer remains current.

8. Enrichment Data Source Ecosystem

TTF cross-references every analyzed report against 10 curated enrichment data sources, each automatically refreshed via weekly CI/CD pipelines:

1. **LOLDrivers:** Complete catalog of vulnerable and malicious kernel drivers with hashes, CVEs, and publisher metadata. Flags techniques involving kernel-level compromise paths [7].
2. **DeTT&CT:** Technique-to-data-source visibility mapping. Identifies exactly which telemetry sources cover which ATT&CK techniques, revealing detection blind spots.
3. **Elastic Detection Rules:** 1,000+ open-source detection rules parsed with MITRE technique mappings, severity levels, and risk scores. Enables instant detection gap analysis [8].
4. **MITRE ATT&CK for ICS:** Full Industrial Control Systems threat matrix. Extends all existing analytics to OT/ICS environments without separate tooling [9].
5. **RE&CT Framework:** 45+ incident response actions mapped to ATT&CK techniques across 6 response stages. Ensures that every finding has a defined response playbook.
6. **MITRE ATLAS:** 40+ adversarial techniques targeting AI/ML systems, including prompt injection, model poisoning, and LLM jailbreak vectors.
7. **VERIS-ATT&CK Mappings:** Real-world incident classification patterns (from the Verizon DBIR methodology) mapped to ATT&CK for compliance benchmarking.
8. **OSSEM:** 40+ security event sources across Windows, Linux, and Cloud with ATT&CK technique coverage metadata. Maps your log sources to detection capabilities.
9. **OpenSSF Scorecard:** Supply chain security scores for critical open-source packages. Tracks code review practices, maintenance activity, and known vulnerabilities.
10. **MITRE Engage:** 26+ adversary engagement activities for deception, denial, and disruption operations mapped to ATT&CK techniques.

These sources are in addition to the runtime threat intelligence feeds (Shodan, HIBP, ThreatFox, URLhaus, Ransomwatch, CISA KEV) that are refreshed at every application startup. Combined, TTF maintains one of the broadest enrichment ecosystems available in a single platform.

9. Novel Analytics Suite

Beyond standard compliance mapping and risk scoring, TTF includes 13 purpose-built analytics engines that provide intelligence unavailable from any other single platform:

1. **Technique Maturity Index (TMI):** Per-technique 0-100 readiness score across 6 dimensions -

Detection, Compliance, Testing, Remediation, Intelligence, and Assessment.
2. **Defensive Debt Score:** Compound risk metric calculated from overdue remediation tasks, unvalidated detections, recurring techniques, and compliance gaps. Quantifies accumulated security neglect.
3. **Compliance Blast Radius:** Cross-framework regulatory impact analysis showing which controls are violated per technique, with estimated penalty exposure across HIPAA, PCI-DSS, GDPR, SEC, FISMA, NERC CIP, NYDFS, and CCPA.
4. **Kill Chain Velocity Tracker:** Measures how quickly attackers progress through kill chain phases. Detects acceleration or deceleration trends across reports.
5. **Attacker ROI Calculator:** Adversary economic modeling - effort, cost, success probability, and potential payout per attack path. Identifies which techniques offer attackers the best return.
6. **Adversary Adaptation Predictor:** Game-theoretic forecast of how adversaries will shift tactics based on your defensive improvements.
7. **Industry Threat Weather Map:** Anonymized peer data showing active, intensifying, or declining threats across your sector.
8. **Entity Persistence Map:** Cross-report infrastructure graph revealing repeatedly targeted assets and chronic compromises.
9. **Technique Dependency Graph:** Causal directed graph showing attack supply chains - which techniques enable which, identifying single points of failure in the kill chain.
10. **Finding Correlation Fingerprints:** Root-cause clustering via CWE/OWASP similarity to group reports by underlying security problems rather than surface-level symptoms.
11. **Threat Actor Convergence Warnings:** Detects when multiple unrelated threat actors show simultaneous interest in the same techniques or CVEs - an early warning signal.
12. **Shadow IT Discovery Engine:** Identifies unmanaged infrastructure by mining entity graphs from offensive reports.
13. **Predictive Compliance Drift:** Detects where emerging threats are concentrating in areas of weak compliance coverage before a regulatory gap becomes an incident.

Each analytics engine is deterministic - no LLM is used for analytics computation. Results are derived from SQL aggregation, statistical modeling, and graph analysis, ensuring reproducibility and audit readiness.

10. The Command Center & Executive Operations

Security programs fail in the seams between tools. The CISO Command Center closes those seams by synthesizing every TTF module, including reports, vulnerabilities, detections, incidents, regulatory obligations, and analytics, into a single posture score with ranked next actions.

- **Unified Posture Score:** A daily-snapshotted, trend-tracked composite across every surface, with a 30-day sparkline and per-module drill-downs answering "why did this move?"
- **Ranked Next Actions:** The platform continuously ranks the highest-leverage actions across all modules - the one prioritized to-do list for the program.
- **Board-Ready PDF:** A one-click executive summary rendering current posture, trends, and priorities for board packets.
- **Morning Brief:** An auto-generated daily stand-up of what changed overnight, including new incidents, KEV additions, detection results, and approaching deadlines.
- **Threat Digest & On-Call Brief:** Periodic org-specific intelligence digests, plus handoff-ready

situational summaries for the incoming on-call.

The Virtual-CISO Toolkit

TTF also packages the recurring executive workload of a security leader:

- **Budget Memo Composer:** Turns a security ask into a finance-grade memo backed by the same Monte Carlo engine that powers the Rhino Score, including expected-loss modeling, risk-reduction assumptions, and peer-spend comparison, with linked findings as evidence.
- **Board & Meeting Prep:** Paste an agenda and receive a briefing pack tuned to the audience, whether that is the board, the executive team, finance, the audit committee, or a regulator.
- **Security Program Maturity:** Continuous maturity scoring computed from what the organization actually does in the platform, not a self-assessment survey.
- **Security Roadmap Sequencer:** Orders the backlog of security work by risk reduction per unit effort, with dependency-aware sequencing.
- **Peer CISO Benchmark:** Posture, coverage, and velocity compared against anonymized industry and size-cohort peers.
- **Cyber Insurance Gap Analyzer:** Maps evidenced posture against common policy warranties and exclusions.
- **M&A Due Diligence:** Rapid security due-diligence memos for acquisition targets, exportable as Markdown or PDF.
- **Policy Drift Detector:** Flags where written security policy has drifted from observed practice in the organization's own data.
- **Analyst Skill-Gap Analysis**: Identifies the technique and tooling areas where the team's demonstrated coverage is thinnest, with a training plan.

11. Incident Response & the Regulatory Engine

Breach notification law has become a minefield of overlapping deadlines: GDPR's 72 hours, the SEC's four business days, HIPAA, state attorneys general, and sector-specific regimes. Missing one clock converts a security incident into a legal one.

When an incident is declared in TTF, the regulatory engine automatically evaluates the incident's characteristics (data types, jurisdictions, record counts, industry, and public-company status) against a maintained catalog of notification regulations, determines which obligations apply, computes each deadline from the correct legal basis (discovery, confirmation, or materiality determination), and starts the clock.

- **Live Regulatory Countdown:** A persistent header indicator on every page shows the tightest active clock, shifting from green to amber to pulsing red as deadlines approach. Deadlines recompute automatically as incident facts are corrected.
- **Incident War Room:** A single working surface per incident that includes timeline, tasks, evidence, obligations, and communications.
- **AI-Drafted Notifications:** Regulator and customer notifications drafted from templates keyed to each specific regulation, then reviewed, versioned, and delivery-tracked with content hashing.
- **Governed Waivers:** Waiving an obligation requires a signed-off entry in the Decision Journal, an auditable record of who accepted which risk, why, and when it expires.
- **Post-Mortem Generator:** Structured after-action reports composed from the incident record:

timeline, root cause, action items, and lessons learned.

12. Detection Engineering & Alert Triage

Detection Engineering Advisor

From entered findings or selected techniques, the advisor produces a prioritized detection plan: which techniques to detect first (weighted by exposure, adversary interest, and existing coverage), which rules to deploy, and per-verdict supporting evidence. Recommendations draw on the Sigma and Elastic rule corpora plus DeTT&CT visibility mappings.

Detection Stack Translation

Organizations declare their stack once - Splunk, Microsoft Sentinel, Elastic, CrowdStrike, QRadar, Chronicle, and 30+ other platforms are supported - and every Sigma rule and detection recommendation is translated into the native query language of that stack.

Alert Triage Intelligence

SIEMs can forward alerts to a signed TTF webhook, where each alert is re-scored against the organization's actual exposure: its open CVEs, observed techniques from offensive reports, matched threat-actor interest, and KEV status. Analyst feedback (confirm/dismiss) tunes future scoring. The result: the alert queue is re-ordered by what can actually hurt this organization, not by vendor-assigned severity.

Detection Rule Validation

A dedicated validation queue tracks whether deployed rules have actually been tested (pending, queued, validating, completed), and unvalidated detections feed the Defensive Debt score where untested detection is quantified as risk, not assumed as coverage.

Collective Intelligence

Privacy-first, anonymized cross-tenant signal sharing shows which techniques and CVEs are trending across the TTF fleet, giving every tenant early warning from the community's aggregate exposure. Participation is controllable with a simple opt-out.

13. Vulnerability Management & KEV Intelligence

TTF includes a full vulnerability lifecycle manager that unifies findings from offensive reports and scanner imports (Burp XML, Nessus CSV, Qualys XML) into a single tracked queue.

- **SLA Enforcement:** Every vulnerability carries a due date derived from severity-based SLA policy. CISA KEV-listed vulnerabilities automatically receive accelerated due dates aligned with BOD 22-01. Breach lists and MTTR metrics make SLA performance visible and reportable.
- **Newly-KEV'd Alerts:** When CISA adds a CVE that already appears in the organization's reports, affected findings are flagged and the right owners are notified automatically.
- **KEV Intelligence Suite:** KEV burndown tracking, remediation velocity, vendor concentration analysis, ransomware exposure correlation, a detection heatmap, per-CVE adversary emulation

plans, and a predictive model estimating which CVEs are most likely to be added to KEV next.

14. Differentiators

Feature Domain	Implementation & Technological Edge	Strategic Value to the CISO
Advanced AI Pipeline	Google Gemini 3 Pro (1M+ Token Context) paired with a deterministic regex engine for high-fidelity T-Code extraction	Eliminates manual report triage; converts "dead data" in PDF format into actionable defensive intelligence with near-zero latency
Financial Risk Modeling	Monte Carlo Risk Engine (10,000 iterations) translating technical debt into "Expected Financial Loss" based on an average $56k IR baseline	Quantifies negligence and provides a defensible, FICO-style Rhino Score (0-850) for board-level reporting
Predictive Threat Modeling	ML-Powered Forecasting of technique frequency, APT targeting, and probabilistic risk trajectories	Shifts posture from reactive defense to proactive mitigation by anticipating the adversary's next move
Framework Orchestration	Deterministic multi-framework cross-walking across PCI-DSS v4.0, ISO 27001, CIS v8, NIST 800-53, and D3FEND	Automates audit evidence generation, reducing annual compliance preparation time by an estimated 75%
Unified Intelligence	10 Aggregated Threat Feeds (Shodan, HIBP, etc.) with Smart ROI Caching (90% API cost reduction)	Provides real-time context on IOCs while maintaining local data persistence for historical trend analysis
Conversational Analysis	RAG-Powered Unified AI Assistant supporting complex natural language (NL) queries across the entire report history	Democratizes data access; allows security leadership to ask "What findings impact our payment gateway?" with immediate answers
Active Validation	Purple Team Exercise Platform with Atomic Red Team integration and asynchronous detection validation	Validates that security controls actually work against real-world APT techniques like APT29 and Lazarus
Advanced Analytics	Custom Widgetized Dashboards and Scheduled Executive Reports delivered via Email, Webhook, or Cloud Storage	Delivers curated, automated reporting to the board, ensuring transparency and continuous ROI visibility
Enterprise Governance	SAML 2.0/OIDC SSO, Immutable Audit Logging, and SOC 2-compliant Data Retention/Archival policies	Ensures enterprise-grade security, multi-tenant isolation, and complete regulatory compliance for data residency
Supply Chain Insight	SBOM Extraction correlating report-identified software dependencies with live CVE/KEV vulnerability databases	Proactively identifies systemic risk within the software supply chain before a vendor-specific breach occurs
Ransomware Intelligence	55 techniques, 11 kill chain phases, 74+ tracked groups, financial impact modeling, and 8 pre-built Purple Team exercises	Provides a single ransomware readiness score with actionable remediation priorities and board-ready financial exposure estimates
Threat Actor	GapMATRIX integration with 290+	Enables "who is targeting us?" analysis -

Feature Domain	Implementation & Technological Edge	Strategic Value to the CISO
Correlation	actor profiles and 300+ CVE-to-actor mappings, weekly auto-sync	attributes findings to specific threat groups and anticipates their next moves
Enrichment Ecosystem	10 curated data sources (LOLDrivers, DeTT&CT, Elastic Rules, ATT&CK ICS, RE&CT, ATLAS, VERIS, OSSEM, OpenSSF, Engage) refreshed weekly	Provides the broadest single-platform enrichment coverage, eliminating the need for 5-10 standalone intelligence subscriptions
Novel Analytics	13 purpose-built engines including TMI, Defensive Debt, Kill Chain Velocity, Attacker ROI, and Adversary Adaptation Predictor	Delivers intelligence unavailable from any other platform - quantifies security neglect, predicts adversary pivots, and measures analyst effectiveness
Critical Findings Alerts	Compound-evidence model evaluating 6 independent signal types; fires only when 2+ signals converge	Eliminates alert fatigue by ensuring only genuinely impactful findings trigger notifications, with tailored remediation recommendations
CISO Executive Dashboard	Real-time KPIs, Rhino Score trends, ransomware exposure, regulatory impact across 8 frameworks, financial loss projections	Provides the Board of Directors with a single pane of glass for security posture, replacing quarterly slide decks with live data
Command Center	Unified posture score with daily snapshots, ranked cross-module actions, 30-day trends, and one-click board PDF	Replaces quarterly slide decks with a live, defensible view of the whole program
Regulatory Response Engine	Automated evaluation of breach-notification obligations (GDPR, SEC, HIPAA, state AGs) with per-obligation deadline computation and live countdowns	Converts regulatory chaos into a managed checklist; deadlines are never discovered after they pass
Virtual-CISO Toolkit	Monte Carlo budget memos, audience-tuned meeting prep, program maturity scoring, roadmap sequencing, peer benchmarking, insurance gap analysis, M&A due diligence	Packages the recurring executive workload of a security leader into evidence-backed, generated deliverables
Detection Engineering	Prioritized detection advisor, stack-native Sigma translation (30+ platforms), and SIEM alert re-scoring against actual organizational exposure	Detection investment goes where the exposure is; the alert queue is ordered by what can actually hurt you
Vulnerability Lifecycle	SLA-enforced tracking with KEV-aware due dates (BOD 22-01), scanner imports, MTTR and breach reporting	Turns vulnerability management from a spreadsheet into an accountable, deadline-driven process

Table 1: Core platform capabilities

15. Security & Trust

TTF is designed for **SOC 2 Type II** readiness. Security is not a feature; it is the foundation.

- **Data Isolation:** Strict multi-tenancy is enforced at the ORM layer via `organization_id` checks on every query.

- **Auditability:** Comprehensive `AuditLog` captures every state-changing operation (POST/PUT/DELETE) and sensitive read, recording User, IP, Resource, and Action [10].
- **Access Control:** Granular RBAC (Viewer/Editor/Admin) combined with Enterprise SSO (SAML 2.0 / OIDC) ensures least-privilege access.
- **Resilience:** Rate limiting middleware (`RateLimitUsage`) protects against DoS and abuse, configurable per IP, User, or Organization.
- **Global Data Sovereignty and Residency:** To meet the stringent requirements of global enterprises and regulatory bodies (GDPR/CCPA), TTF supports Multi-Region Deployment. Organizations can explicitly define geographic data residency rules (e.g., US, EU, APAC) to ensure that sensitive offensive intelligence remains within specific jurisdictional boundaries. Combined with SOC 2-compliant automated archival and data retention policies, TTF provides the governance foundation required for multinational security operations.
- **Edge Protection:** All API traffic passes through a global external load balancer with Cloud Armor with per-IP rate-limiting bans and WAF rulesets (SQLi, XSS, LFI, RCE, scanner detection) before reaching the application. Direct access to the application runtime is disabled.
- **Webhook Integrity:** Inbound webhooks (billing, SIEM alerts) require cryptographic signature verification, with an idempotency ledger preventing replay and out-of-order processing.
- **Session Security:** HMAC-signed session tokens with server-side revocation, meaning sessions can be terminated centrally, not merely expired client-side.
- **Continuous Assurance:** The platform undergoes recurring full-surface code audits (security, logic, and UX), with findings tracked to closure and verified by an automated post-deploy end-to-end test gate against production.

Control Domain	Implementation Mechanism	Assurance Objectives
Tenant Isolation	Mandatory `organization_id` filtering at ORM/SQL level	Prevents cross-tenant data leakage
Auditability	Immutable `AuditLog` records User, IP, Resource, Action, Timestamp	Supports SOC 2 & forensic investigations
Identity & Access	Enterprise SSO (SAML/OIDC) with granular RBAC	Enforces Least Privilege & centralizes de-provisioning
Data Retention	Configurable archival (Cloud Storage) or hard deletion policies	Complies with GDPR/CCPA data minimization

Table 2: Platform security & governance controls

16. Business Value & ROI

Assumption: Standard Enterprise Audit costs $150k/yr; a fully loaded Senior Security Analyst costs $180k/yr..

Operational Area	Efficiency Gain	Est. Annual Savings
Audit Preparation	Automated evidence generation across 11+ frameworks reduces prep time by 75%	$40,000+
Threat Intel Costs	10 enrichment sources + smart caching replaces 5-10 standalone subscriptions	$85,000+
Analyst Productivity	Automated report ingestion, mapping, and analytics (saves	$35,000+

Operational Area	Efficiency Gain	Est. Annual Savings
	~10 hrs/report)	
Ransomware Preparedness	Continuous readiness scoring replaces annual tabletop exercises	$25,000+
Incident Response	Critical findings alerts reduce mean time to detect compound threats	$15,000+
Total Estimated Annual Savings		**$200,000+**

Table 3: Cost savings assumptions

Illustrative Scenario: A 48-Hour Transformation

The following is a modeled scenario with hypothetical figures, illustrating the intended deployment pattern. Consider a mid-sized financial institution with a historical repository of 12 penetration test reports spanning three years and a Recurring TTP Rate of 22%, repeatedly paying for the discovery of identical vulnerabilities without successful remediation.

Within 48 hours of deploying the Threat Teaming Framework:

- Automated Ingestion recovered 400+ hours of manual analyst correlation time by instantly parsing and indexing the entire historical report backlog.
- Predictive Modeling identified a high probability of future targeting by APT29 based on emerging TTP trends identified across the historical data and live threat feeds.
- Mobilization pushed 15 high-priority remediation tasks directly into Jira, each deterministically mapped to specific CIS v8 controls required for an upcoming regulatory audit.

The modeled outcome: a security posture improvement from a Rhino Score of 540 to 710 within a single quarter - quantifiable, defensible evidence of defensive maturity for the Board of Directors.

17. Implementation & Measurable KPIs

Deployment: Containerized (Docker) deployment to Cloud Run. Time to Value: <24 Hours (SaaS).

Success Metrics (KPIs)

1. **Rhino Score Improvement:** Month-over-month increase in aggregate security score (Target: >700).
2. **Detection Coverage Delta:** Net new techniques covered by validated detection rules per quarter.
3. **Recurring TTP Rate:** Percentage of findings in new reports that match previously identified techniques (Target: <5%).
4. **Remediation Velocity / MTTR:** Average days to close Critical findings and SLA-tracked vulnerabilities, with KEV items measured against BOD 22-01 timelines.
5. **Regulatory Deadline Adherence:** Percentage of notification obligations completed inside their legal window (Target: 100%).
6. **Posture Trend:** Direction and slope of the Command Center posture score over trailing 30/90 days.

7. **Audit Evidence Readiness:** Time to generate a control coverage report for a specific framework (Target: <5 minutes).

18. Conclusion

Passive security is a liability. The Threat Teaming Framework transforms the necessary expense of offensive security into a strategic asset. By operationalizing data that currently sits dormant in the "PDF Black Hole," TTF provides the CISO with defensible metrics and ransomware readiness intelligence, the engineer with actionable context enriched by 10 data sources and 290+ threat actor profiles, and the organization with measurable resilience backed by 13 purpose-built engines.

The platform has evolved from a report ingestion tool into a complete security operations and governance layer: a Continuous Threat Exposure Management engine that tracks 74+ ransomware groups and the broadest single-platform enrichment ecosystem available; an incident response system that computes and counts down every regulatory notification deadline; a detection engineering advisor that translates recommendations into the native language of 30+ SIEM and EDR platforms; a vulnerability manager with KEV-aware SLA enforcement; and an executive toolkit that turns all of it into board-ready deliverables that are synthesized daily into a single Command Center posture score.

Next Steps: Close the Gap

The window of exposure is closing. Do not let your high-value offensive intelligence sit dormant in a "PDF Black Hole." Transform your security posture from a reactive liability into a strategic asset today.

- Schedule a Strategy Session: hello@hoffmann.holdings
- View a Live Demo: https://threatteamingframework.com/demo
- Start Your Posture Transformation: Deploy in under 24 hours via our containerized SaaS delivery model.

---

19. References

[1] *Hybrid AI-Deterministic Extraction Pipeline Specifications (internal specification, available under NDA)*
[2] *Proprietary Monte Carlo Financial Risk Engine Documentation (internal specification, available under NDA)*
[3] *Autonomous Security Control Enrichment Engine (11+ Frameworks) (internal specification, available under NDA)*
[4] *Multi-Source Threat Intelligence Aggregation Methodology (internal specification, available under NDA)*
[5] *SigmaHQ Detection Rules - github.com/SigmaHQ/sigma*
[6] *GapMATRIX Community Threat Actor Intelligence - github.com/EssexRich/ThreatActors-TTPs*
[7] *LOLDrivers Vulnerable Driver Catalog - loldrivers.io*
[8] *Elastic Detection Rules - github.com/elastic/detection-rules*
[9] *MITRE ATT&CK for ICS - attack.mitre.org/techniques/ics*
[10] *Relational Architecture for GRC Framework Mapping (internal specification, available under NDA)*