Threat Teaming Framework: Operationalizing Offensive Intelligence

Closing the Gap Between Detection and Remediation via Automated Threat Modeling

Publish Date: April 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) 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.

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.

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.

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 with real-time threat intelligence, maps them to GRC frameworks, and calculates financial risk to drive prioritized remediation.

The TTF is designed to operationalize the five stages of the Gartner Continuous Threat Exposure Management (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 3,800+ 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.

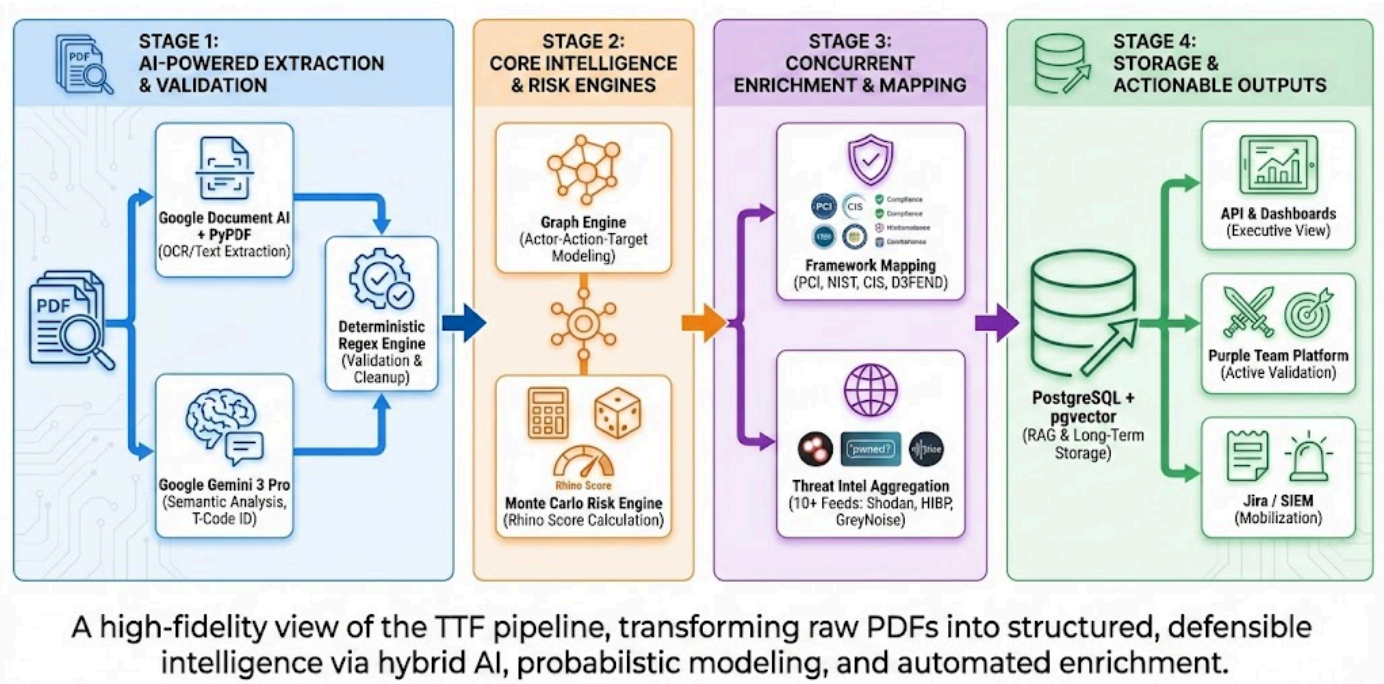

Figure 1: The ingestion pipeline

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 3.13 (FastAPI).
- **Database:** PostgreSQL 14+ with pgvector and JSONB for flexible schema handling.
- **AI/ML:** Google Gemini 3 Pro Preview via google.generativeai SDK.
- **Asynchronous Task Queue:** Cloud Tasks for long-running validation jobs.
- **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.

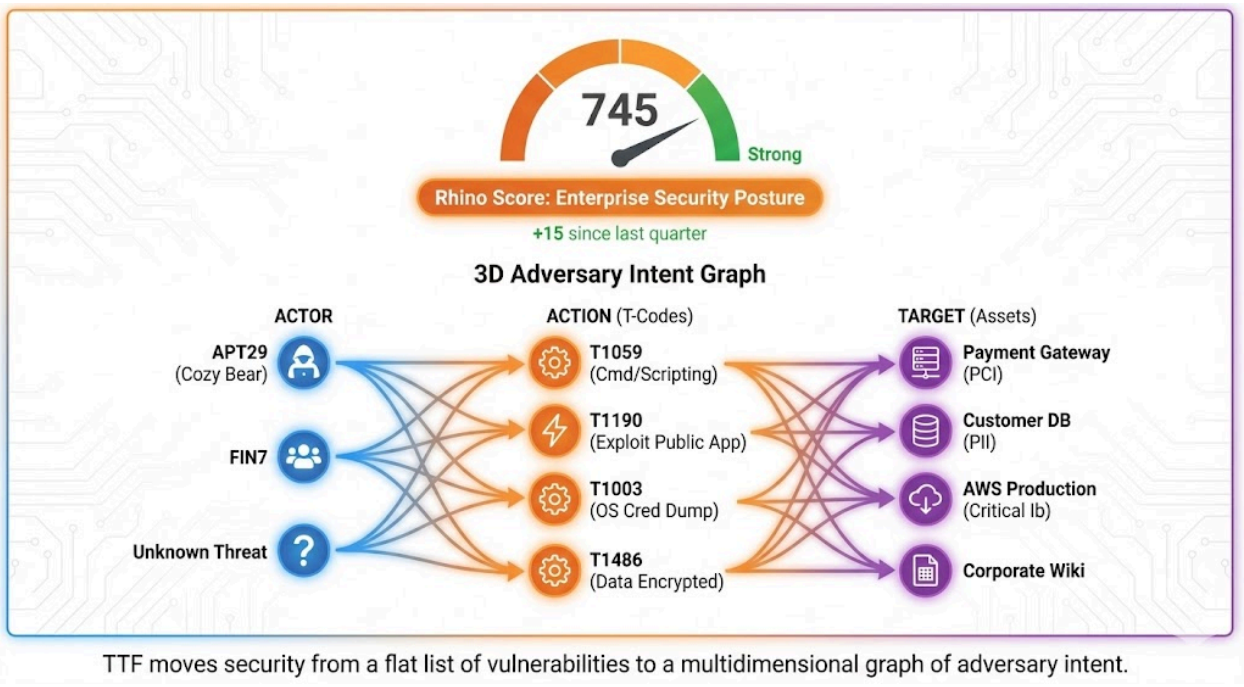

Figure 2: Rhino score based on Monte Carlo simulations using multiple variables

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. Asynchronous Detection Validation

The platform includes a dedicated validation queue. When a new detection rule is proposed, it is queued for asynchronous validation. The system tracks the status (`pending`, `queued`, `validating`, `completed`) and stores the result, ensuring that the UI remains responsive while heavy-lifting verification occurs in the background.

6. Differentiators

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 1: Platform security & governance controls

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, GreyNoise, 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	Validates that security controls actually

Feature Domain	Implementation & Technological Edge	Strategic Value to the CISO
	Atomic Red Team integration (3,800+ tests) and asynchronous detection validation	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 SOC2-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

Table 2: Core platform capabilities

7. 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 [5].
- **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 SOC2-compliant automated archival and data retention policies, TTF provides the governance foundation required for multinational security operations.

8. Business Value & ROI

Assumption: Standard Enterprise Audit costs $150k/yr; Senior Security Analyst loads cost $180k/yr.

Operational Area	Efficiency Gain	Est. Annual Savings
Audit Preparation	Automated evidence generation reduces prep time by 75%	$40,000+
Threat Intel Costs	Consolidated aggregation & caching (vs. 5+ standalone subs)	$85,000+
Analyst Productivity	Automated report ingestion & mapping (saves ~10 hrs/report)	$35,000+
Total Hard Savings		**$160,000+**

Table 3: Cost savings assumptions

The TTF Impact in Action: A 48-Hour Transformation

A mid-sized financial institution possessed a historical repository of 12 PenTest reports spanning three years. Before deploying TTF, their Recurring TTP Rate was 22%, indicating they were 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 Result: The institution's security posture improved from a Rhino Score of 540 to 710 within a single quarter, providing the Board of Directors with quantifiable, defensible evidence of defensive maturity.

9. Implementation & 6 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:** Average time (days) to close "Critical" findings as tracked in the `RemediationTask` table.
5. **Threat Intelligence Hit Rate:** Percentage of report IOCs correlated with active threat feeds (indicates relevance of testing).
6. **Audit Evidence Readiness:** Time to generate a control coverage report for a specific framework (Target: < 5 minutes).

10. 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, TTF provides the CISO with defensible metrics, the engineer with actionable context, and the organization with measurable resilience.

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.

References

[1] *Hybrid AI-Deterministic Extraction Pipeline Specifications*
[2] *Proprietary Monte Carlo Financial Risk Engine Documentation*
[3] *Autonomous Security Control Enrichment Engine*
[4] *Multi-Source Threat Intelligence Aggregation Methodology*
[5] *Relational Architecture for GRC Framework Mapping*