INTEZER

AI-Adapted Security

A Consumable Blueprint for Real Security Outcomes

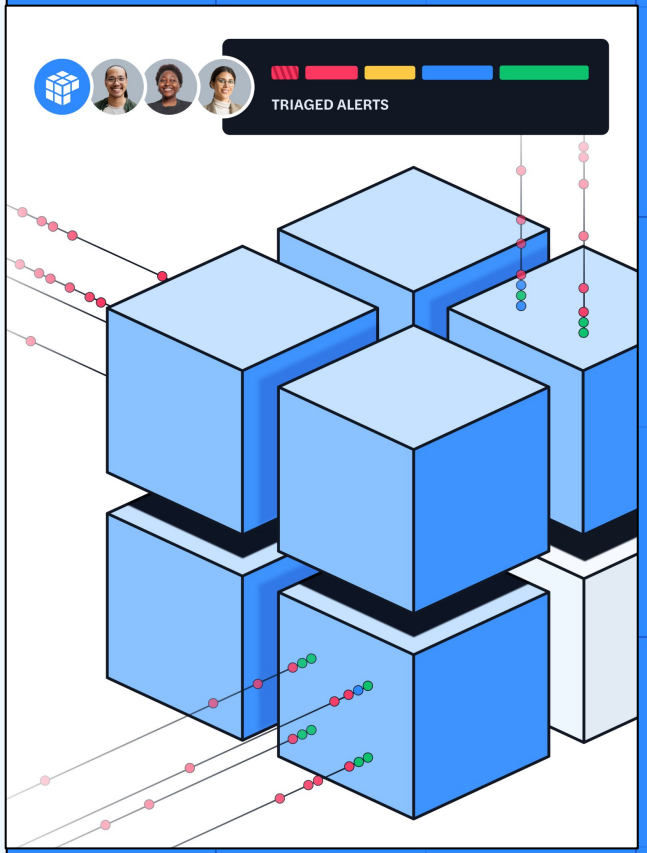

Mitchem Boles - Field CISO

April 9th, 2026

The Security Ops Model Is Broken

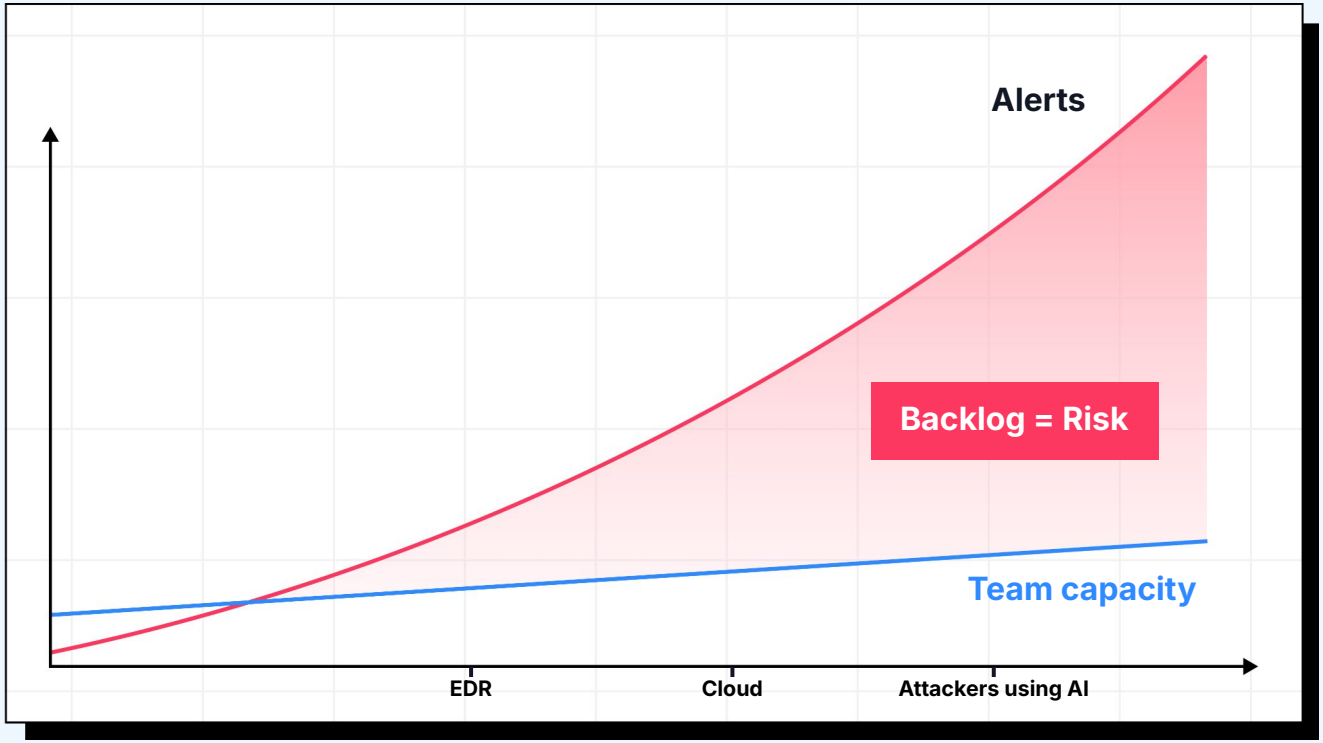

Humans don’t scale as fast as the attack surface gets bigger

Investigation reserved for a small subset of alerts

Depth and accuracy suffer while autonomy lags far behind

Anthropic’s Project Glasswing AND Mythos

WHAT THIS MEANS

1000s

of zero-day vulnerabilities in
every major OS and browser

A single model found bugs that survived **27 years** of expert human review and millions of automated security tests.

It can **autonomously chain exploits** for full system compromise — too dangerous for public release.

What took decades to build can now be broken in minutes.

The response must be equally immediate.

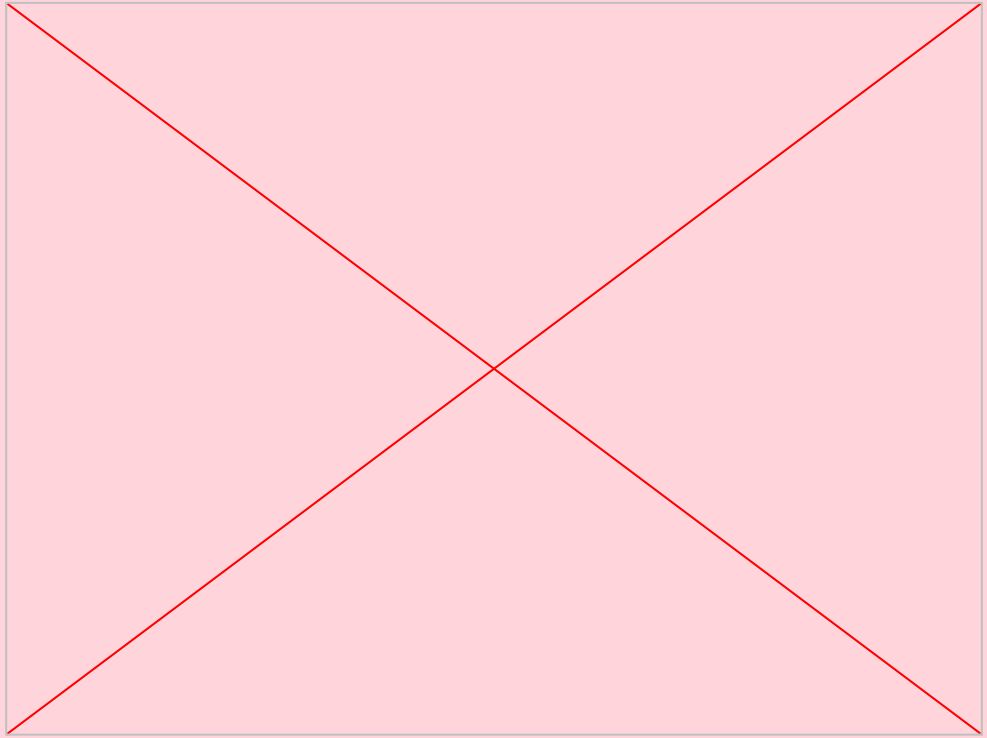

The inflection point → Anthropic's Claude Mythos Preview

an unreleased frontier model too powerful for public access

THINK Operating Model

A reimagined security operations experience that is:

- Proactive
- Scalable
- Outcome-Driven
- Self-improving

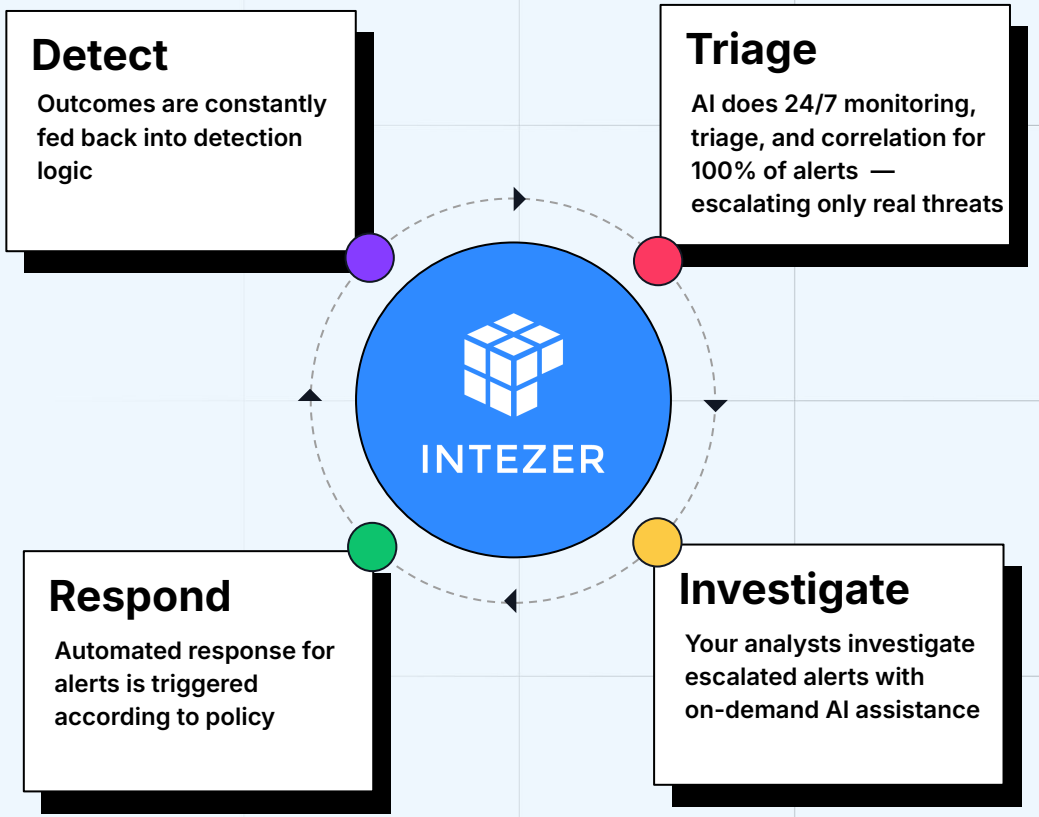

Next-State Operating Model Requirements

Removed Risk Acceptance & Program Maturity

100% of alerts reviewed, triaged, & focused on verdict for minimal escalation.

Spend Optimized - Time Saved

Full coverage across the SecOps alert stack at lower cost with unlimited triage.

Accuracy, Consistency, & AI Transparency

Mitigate genuine threats, full audit trail, & a private instance with AI guardrails.

A Fundamental Shift: From Effort to Outcomes

The old model is broken. Legacy systems, lack of integrations, risk acceptance through reduced investigations, and malicious actors ever adapting. The new vision is a next-state operating model built for machine-speed operations with human oversight.

Machine Speed + Human Clarity = The NEW SOC

How We Get There - 3 Pillars of Transformation

A successful transformation rests on three pillars: **empowering your people**, **automating your processes**, and **integrating intelligent technology**. Intezer is the core that enables this revolution.

People

Elevated & Empowered

- End alert fatigue
- Upskill the team
- Tier-less analyst model with new roles i.e. automation engineer, unified analyst, etc.
- Create high-value roles

Process

Automated & Outcome-Driven

- Move from hours to minutes
- Achieve full alert triage
- Enable proactive, automated remediation
- Optimizing configuration of security tools and ensure best practices

Technology

Intelligent & Integrated

- Tools landscape mapping
- Unified security components
- Core tools provide real, value-driven AI capabilities
- Integrated stacks for easy wins

The ADAPT AI Framework

A strategic roadmap for modernizing Security Operations by aligning People, Process, and Technology.

Assess & Align

Establish a baseline and align with business objectives.

People
- Conduct skills inventory
- Engage stakeholders

Process
- Baseline KPIs (MTTD/R)
- Map critical workflows
- Review compliance

Technology
- Rationalize tools
- Analyze telemetry coverage

Design

Architect the target state and create a roadmap.

People
- Define future roles
- Develop training plan

Process
- Architect playbooks
- Establish data governance

Technology
- Design target architecture
- Create phased roadmap

Automate, Augment, AI

Implement intelligent automation to empower analysts.

People
- Train for human-machine teaming
- Re-allocate analysts to threat hunting

Process
- Implement "Tier 0" triage
- Automate containment

Technology
- Deploy AI for alert prioritization
- Integrate Threat Intel Platforms

Perform

Continuously measure and refine for optimal performance.

People
- Establish analyst feedback loops
- Conduct simulation exercises

Process
- Track performance vs. baselines
- Review playbooks quarterly

Technology
- Build SecOps dashboards
- Monitor AI model efficacy

Transform

Embed change into the culture for sustained adoption.

People
- Launch "Security Champions" program
- Develop role-based training

Process
- Codify new operating model
- Integrate reporting into business reviews

Technology
- Drive full platform adoption
- Decommission legacy systems

Measurable Outcomes

From a Reactive Cost Center to a Proactive & Resilient Business Enabler.

Crawl.Walk.Run with Metrics, Prove the Value.

EXECUTIVE & BOARD METRICS

AVERT: Measuring How Fast We Neutralize Real Business Risk

A Action
V Verified
E Elapsed
R Resolution
T Time

From verified threat to confirmed resolution

*"We track **AVERT** — Action-Verified Resolution Time. It measures from the moment we **verify a real threat** to the moment it's **fully resolved**. It's how we cut through noise and show leadership exactly how quickly we **avert real business risk**."*

< 4 min

Median AVERT with AI SOC — down from 69+ days legacy MTTR

Risk Mitigation Velocity
RMV

20x

How quickly verified threats are neutralized before they become business incidents. Measures speed from confirmed threat to zero exposure.

"We stop threats 20x faster than the industry average."

Autonomous Resolution Rate
ARR

85%

Percentage of threats resolved end-to-end by AI without human intervention. Analysts supervise outcomes, not execute playbooks.

"85% of threats are resolved before an analyst touches them."

Verified Resolution Rate
VRR

97%

Every resolved alert has forensic evidence confirming the threat is neutralized. Not just "closed" — verified closed with proof.

"97% of threats resolved with forensic proof of neutralization."

WHAT THIS MEANS FOR THE BUSINESS

- **Lower breach cost exposure** — Threats contained in minutes, not months
- **Scale without headcount** — AI handles volume; analysts handle final judgment after escalation
- **Audit-ready evidence** — Every resolution backed by forensic proof
- **100% coverage, zero gaps** — Every alert investigated at forensic depth

SOC OPERATIONS — METRICS TRANSFORMATION

From Reactive Metrics to AI-Driven Outcomes

Legacy SOC metrics were designed for a world of sampled coverage and manual triage. AI SOC delivers 100% coverage at forensic depth in under 2 minutes.

LEGACY SOC — BEFORE

Mean Time to Detect (MTTD)
~197 days
Conflates detection + triage. Masks how long real threats sit unexamined.

Mean Time to Respond (MTTR)
~69 days
Measures response initiation, not verified resolution. Doesn't confirm the threat was actually neutralized.

False Positive Rate
~45%
Nearly half of analyst effort wasted on non-threats. Primary driver of burnout.

Alert Coverage
~48%
More than half of alerts never investigated due to volume constraints.

AI SOC

AI SOC — AFTER

Mean Time to Containment (MTTC)
< 4 minutes
From first alert to active containment. Measures what matters: stopping the threat.

Verified Resolution Rate (VRR)
97%+
Percentage of real threats confirmed resolved with forensic evidence — not just "closed."

Analyst Focus Efficiency (AFE)
8.2x
Ratio of analyst time on true positives vs. false positives. AI triages noise; humans handle real threats.

Alert-to-Incident Ratio
200:1 → 8:1
AI correlation compresses raw alerts into verified incidents. 96% noise elimination.

ANALYST IMPACT

73%
Reduction in false positive triage time

100%
Alert coverage (every alert investigated)

5.4 hrs/day
Analyst hours recovered for threat hunting

< 2 min
Forensic-depth investigation per alert

AI Adapted SOC Transformation Roadmap

0-3 MONTHS
Foundation

- ✓ Appoint AI Program Sponsor & Lead Analyst
- ✓ Define outcome-based KPIs & map alert flow
- ✓ Automate a noisy alert source & publish SOC charter

3-9 MONTHS
Acceleration

- ✓ Appoint Automation Engineer & launch AI guild
- ✓ Implement AI SOC platform with verdict engine
- ✓ Achieve >75% auto-triaged alerts & reduce MTTR

9-18 MONTHS
Expansion

- ✓ Formalize new roles & upskill entire team
- ✓ Converge stack into a unified AI SOC Platform
- ✓ Deliver >95% alert coverage & one-click audit reports

18-24 MONTHS
Optimization

- ✓ Redeploy staff to threat hunting & target 2x productivity
- ✓ Introduce autonomous remediation for key scenarios
- ✓ Reduce incident cost >40% & achieve MTTR < 10 min

Accuracy & Depth Unlock Autonomy

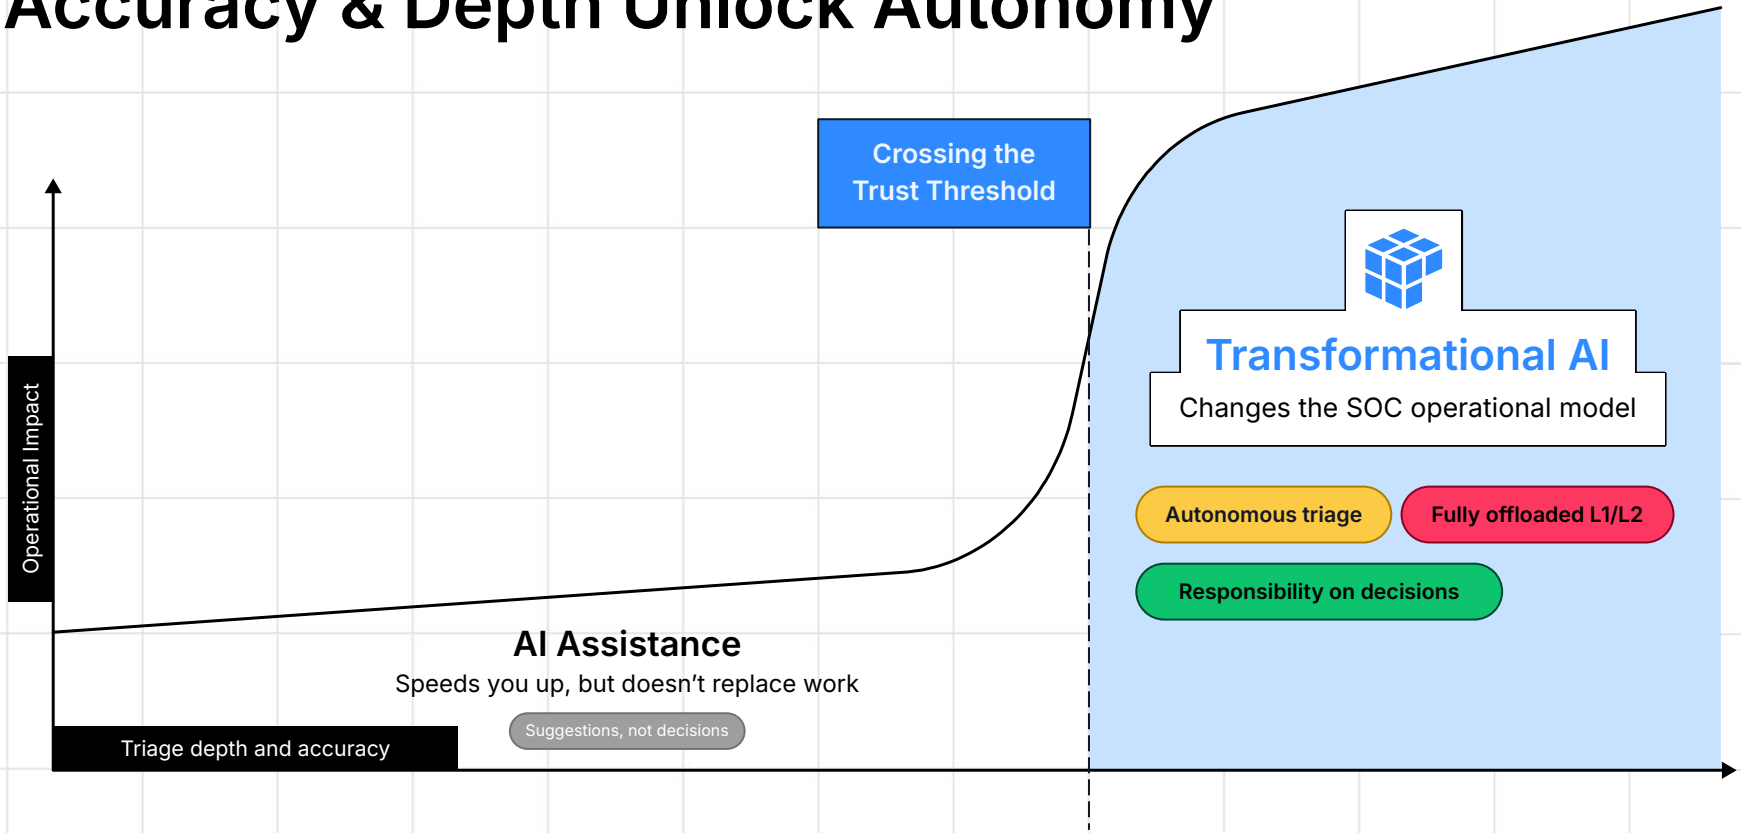

AI SECURITY STRATEGY

The AI-First Enterprise Creates a Fundamentally New Security Problem

AI agents on every endpoint. GPU compute sold as a service. The attack surface isn't growing — **it's changing shape**. Traditional SOC operations can't adapt fast enough without a new model.

HOW THE ATTACK SURFACE CHANGES

Identity & Access — DETECTABLE
Agents act as autonomous identities — API key sprawl, service account overprivilege, multi-tenant auth boundaries

Endpoint — DETECTABLE
Agent processes generate EDR alerts at machine speed — package installs, browser automation, credential access at scale

Data — PARTIAL
Sensitive data flows through prompts and outputs — PII leakage, training data extraction, credentials in AI channels

Cloud & GPU Infrastructure — PARTIAL
Multi-tenant GPU isolation, model serving vulnerabilities, container escape — new infrastructure with new boundaries

Model & Agent Layer NEW — GAP
Prompt injection, RAG poisoning, tool-call hijacking, agent privilege escalation — semantic attacks with no traditional IOCs

WHY THE SOC BREAKS

> *"If an engineer earning $500K is not consuming $250K of AI tokens, that is concerning."*
> — Jensen Huang, March 2026

30K EMPLOYEES × **50+** AGENT ACTIONS/DAY = **1.5M** NEW EVENTS/DAY

On top of existing alert volume — before adding AI-specific detections

40% OF ALERTS ALREADY UNINVESTIGATED

55% OF AI INCIDENTS TRIGGER ZERO SOC ALERTS

82% OF ANALYSTS FEAR MISSING REAL THREATS

The Strategic Insight

This is a **SOC scaling crisis** first and a **new detection domain** second. The alert surge from agent activity will overwhelm operations before AI-specific threats even become the primary concern. The answer is an operating model that triages autonomously at machine scale.

THE PATH FORWARD

Lead AI Security — Don't Wait for the Breach to Build the Playbook

The organizations deploying AI agents at scale are the first to feel the SOC pressure — and the first to solve it. The answer isn't more analysts. It's an operating model built for machine-speed alert volume, combining autonomous triage with AI-specific threat intelligence that your own teams are uniquely positioned to develop.

Absorb the Surge — AI-First SOC Operations

- Deploy **autonomous forensic triage** that investigates 100% of alerts — including the informational and low-severity alerts where real attacks hide
- Reduce human escalation to **~2%** with forensic evidence chains — analysts supervise, they don't sift
- Absorb the 10-50x agent alert surge **without hiring** — the platform scales, your headcount doesn't have to
- Unified pipeline: traditional EDR/identity/cloud alerts **alongside** AI-specific telemetry in one triage workflow

Build What Only You Can — Custom AI Detections

- Your teams know your AI workloads, your agent architectures, your GPU infrastructure — **you build the detections**
- Start with **no-new-tools rules**: agent network egress, credential access patterns, unusual service account behavior, canary tokens in model artifacts
- Layer in **commercial AI-security tools** as the market matures — guardrail verdicts, runtime monitoring, AIDR signals
- Your MDR partner **ingests and operates** your custom detections — you innovate, they run 24/7

Set the Industry Standard — AI Security Leadership

- Publish a **structured AI threat model** covering both agent endpoints and AI cloud — the framework others will adopt
- Map controls across five layers: identity, endpoint, data, cloud/GPU, and model/agent — **prevent, detect, respond**
- Build investigation playbooks for AI-specific scenarios: compromised agent, automated data exfil, rogue model deployment
- Share the approach externally — **the company that secures AI at scale defines the standard** for every enterprise that follows

60-DAY SPRINT TO OPERATIONAL AI SECURITY

THIS WEEK	30 DAYS	60 DAYS	ONGOING
Executive alignment session • MDR evaluation criteria • Data source confirmation • Success metrics defined	AI threat model delivered • Starter detection pack live • Investigation playbooks • Canary tokens deployed	MDR operational with AI telemetry • Custom detection ingestion live • Coverage benchmarks measured	Continuous detection engineering • Quarterly threat model refresh • Emerging AI detection adoption

Schedule Executive Alignment Session
45 minutes to walk the strategy, confirm scope, define success criteria

Confirm Data Sources & Integration Points
EDR, identity, SIEM, cloud, AI agent telemetry, custom detections

Launch the AI Threat Model
2-3 week deliverable covering agent endpoints + AI cloud — assign owners

INTEZER

GET IN TOUCH

Thank You

Questions?

Mitchem Boles - mitchem@intezer.com