ANM
TECH DAY

The Modern SOC: Faster Detection, Smarter Response

Chris Hammer

Cybersecurity Architect

- 25+ Years of IT and Security Experience
- Security, Network, Cloud, and Data Center background
- Built and managed Red / Blue Teams
- I like fast things…If it has an engine, I will drive, ride, or fly it

Key Skills:

- Cybersecurity Risk Management and Program Development
- Aligning security objectives to risk and outcomes
- Validation Services (Red & Blue Team Exercises)
- Enterprise Security Architecture
- Wading through the fluff

What We Will Cover

Agenda

- The Modern SOC: Faster Detection, Smarter Response
- The Attacker Perspective
- What Challenges Does this Cause
- The Defender Evolution
- Real Life Use Case
- Q & A

Zero Trust Principles are Critical

AI Drives the Requirement for Faster Detection and Response

- $10.5T Annual Cost of Cybercrime
- 29 Minutes Average Breakout Time
- 89% Surge in AI-Enabled Attacks YoY

Zero Trust Core Principle's:

Never Trust, Always Verify

Least Privilege Access

ASSUME BREACH

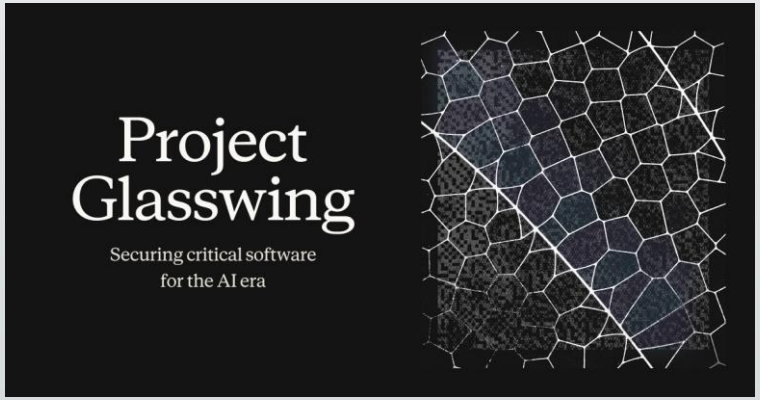

Autonomous AI Hacking?

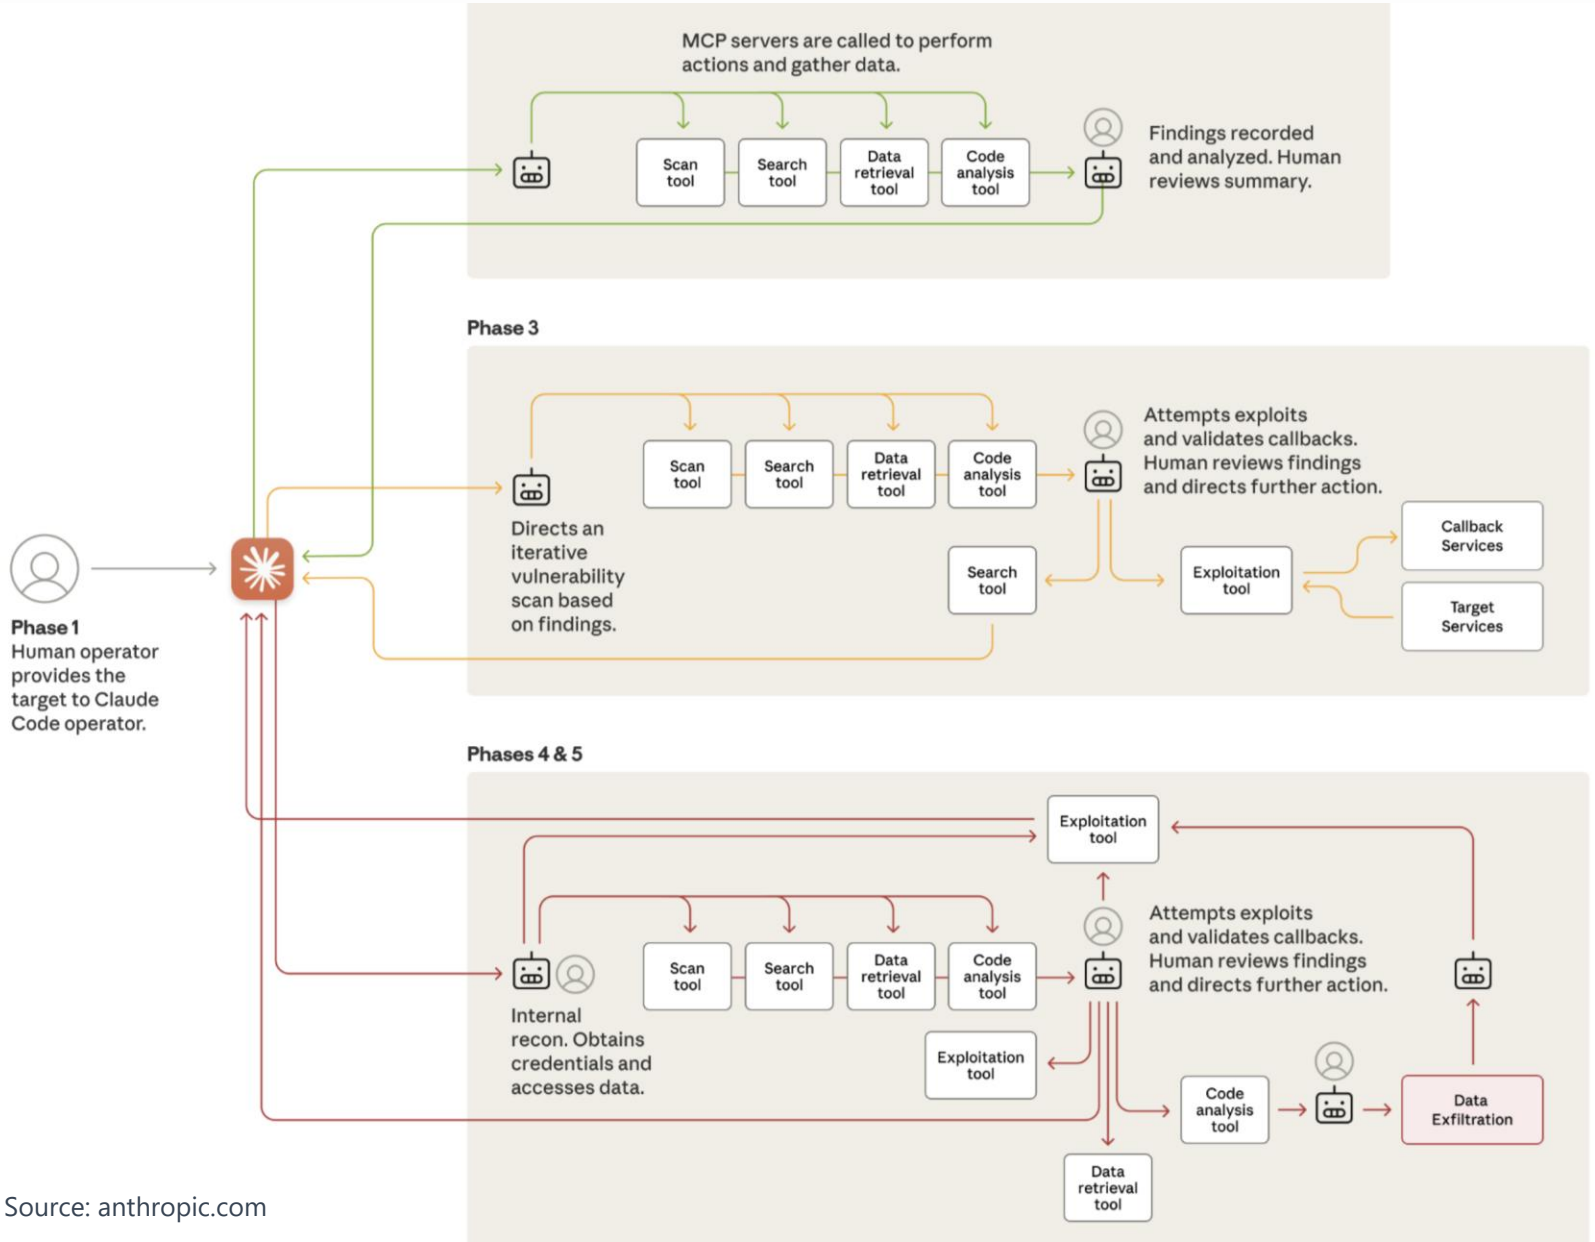

- Target Initialization & Reconnaissance
- Vulnerability Exploitation & Initial Access
- Post-Exploitation & Exfiltration

Source: anthropic.com

How bad is it?

Details in the Noise

Report Claim	Considerations
80% - 90% autonomous / agentic execution	Required custom framework, jailbreaks, human validation
Major inflection point in cyber threats	Hallucinations & unreliability show current limits
Large-scale state-sponsored op	Minimal IOCs, metrics, or independent corroboration shared
AI handled recon -> exploit -> exfil	Techniques echo prior AI-assisted attacks; not a paradigm shift
Proof of agentic weaponization era	AI-Assisted intrusions

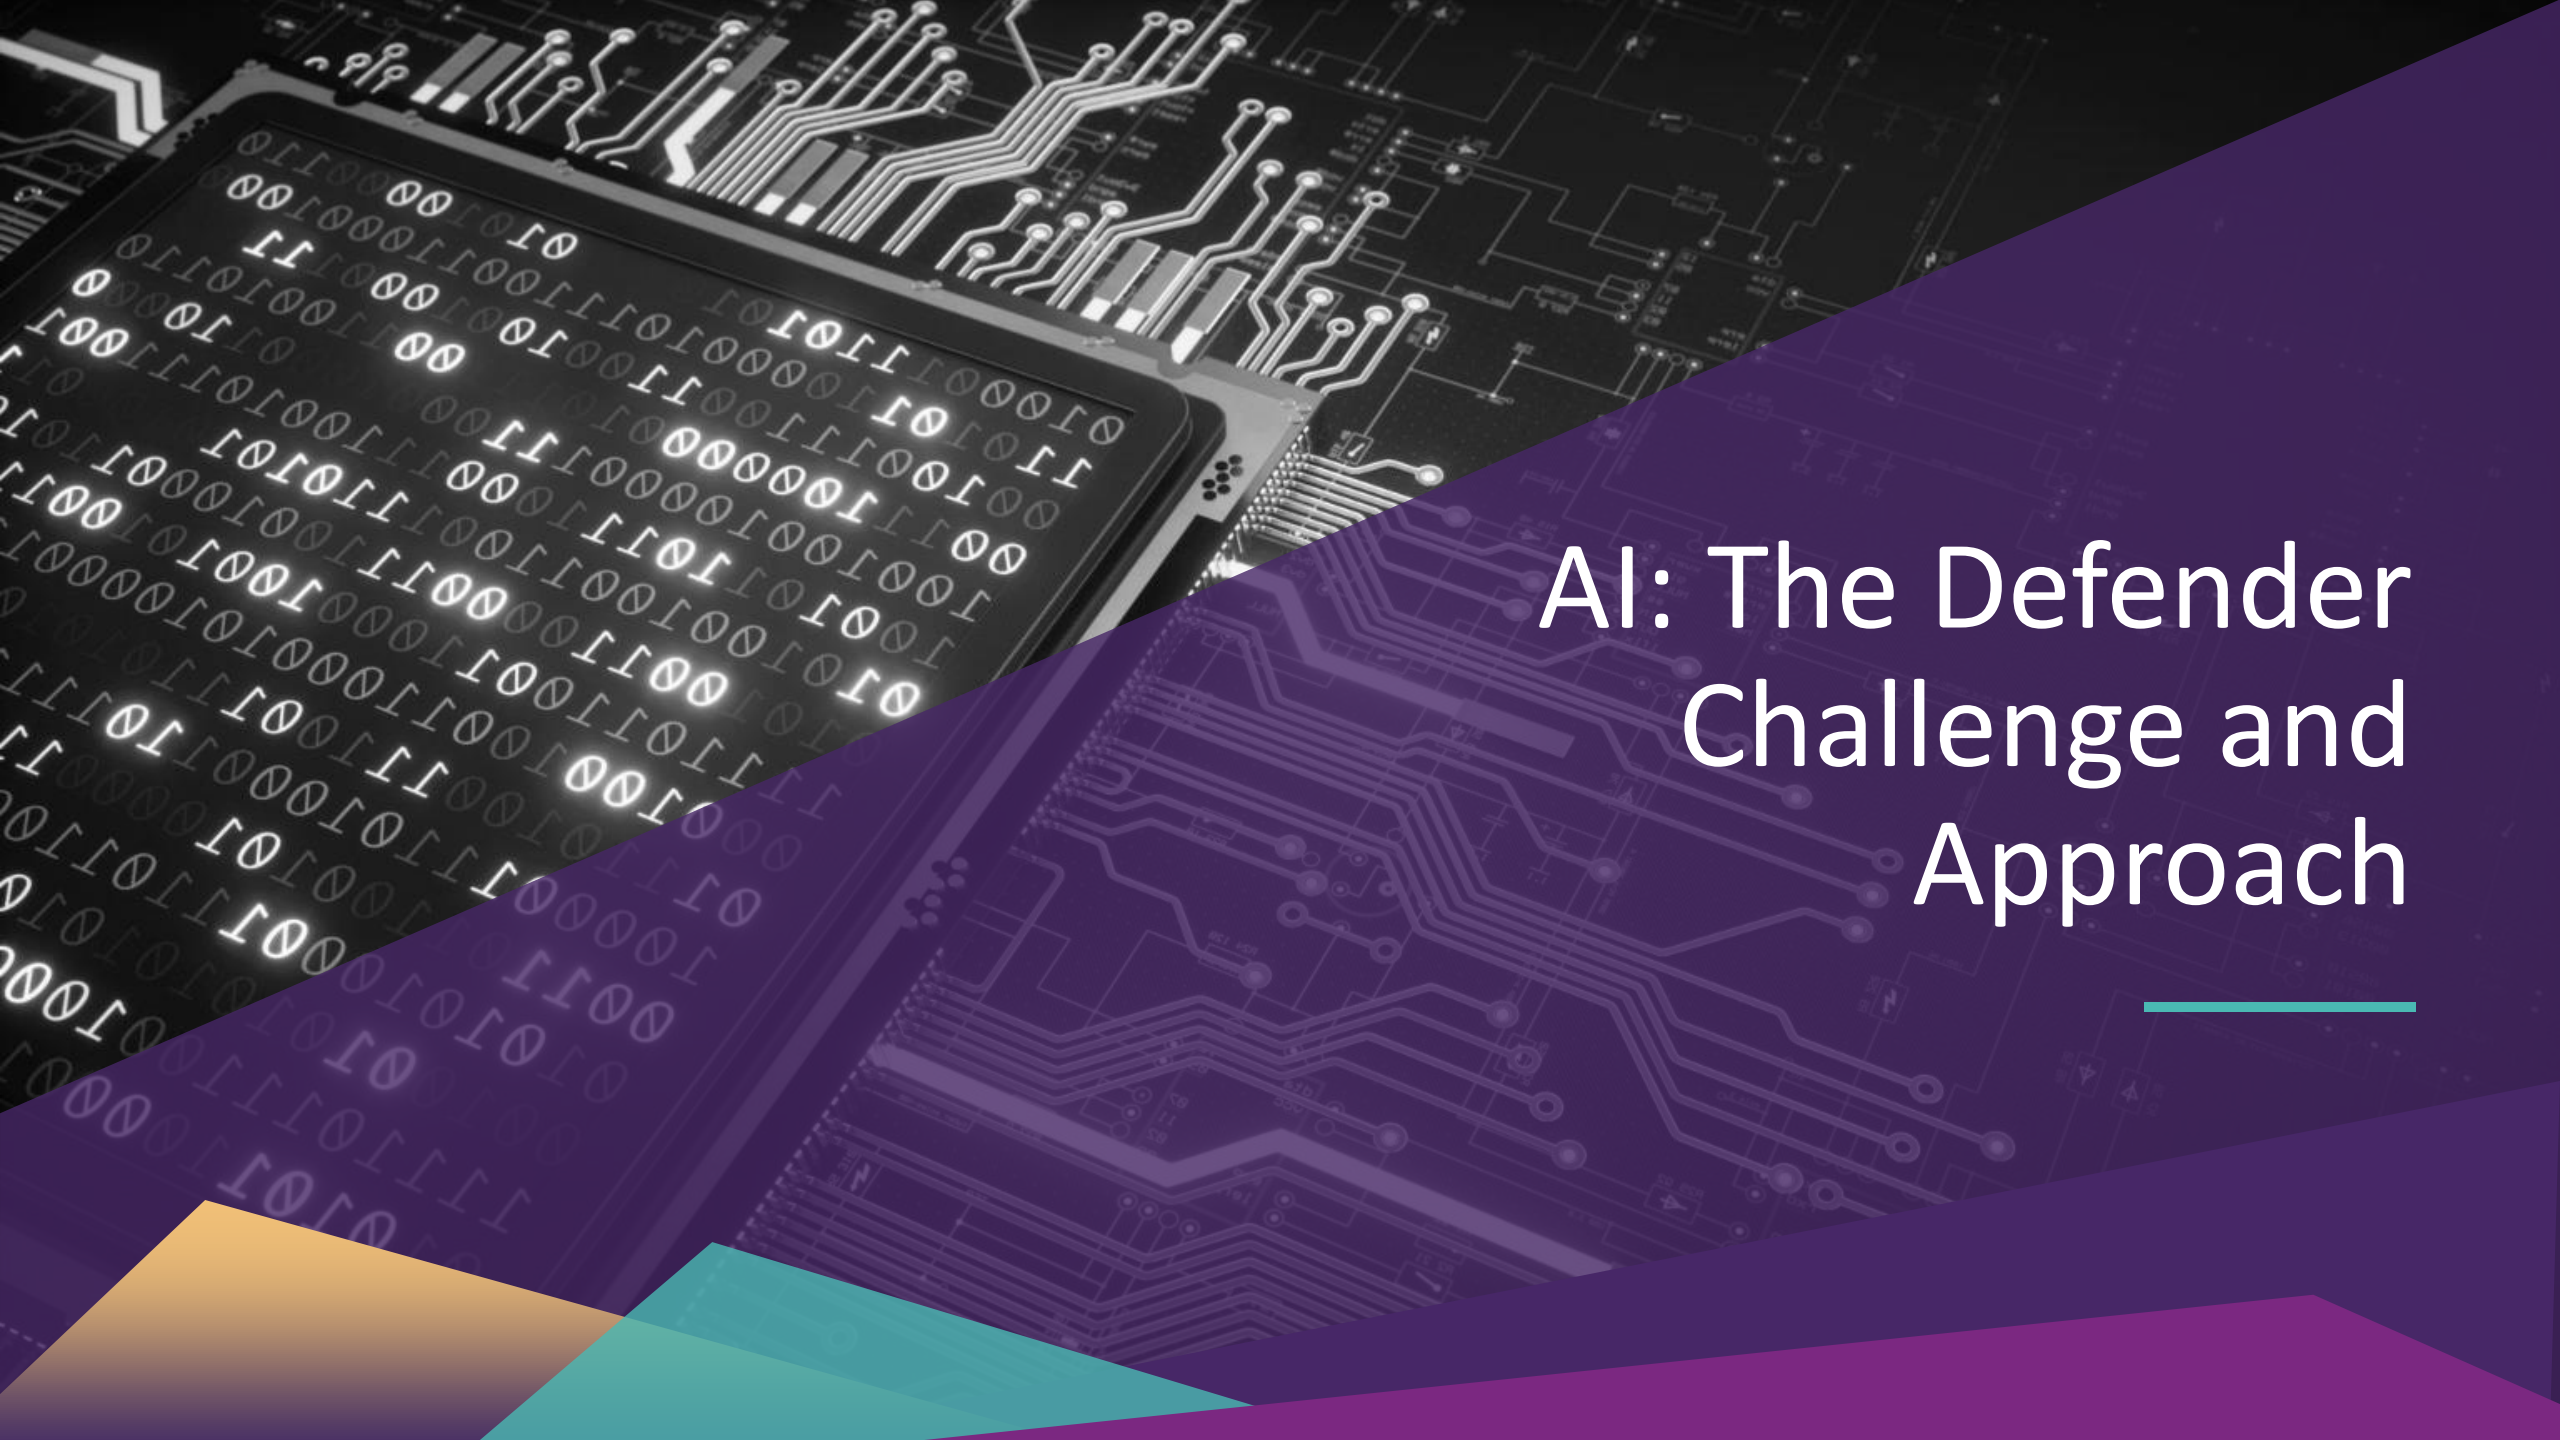

AI: The Defender Challenge and Approach

Critical Zone of Vulnerability

- 70% of SOC analysts experience severe stress
- Up to 64% consider leaving their roles within a year
- Replacing expertise is expensive and challenging

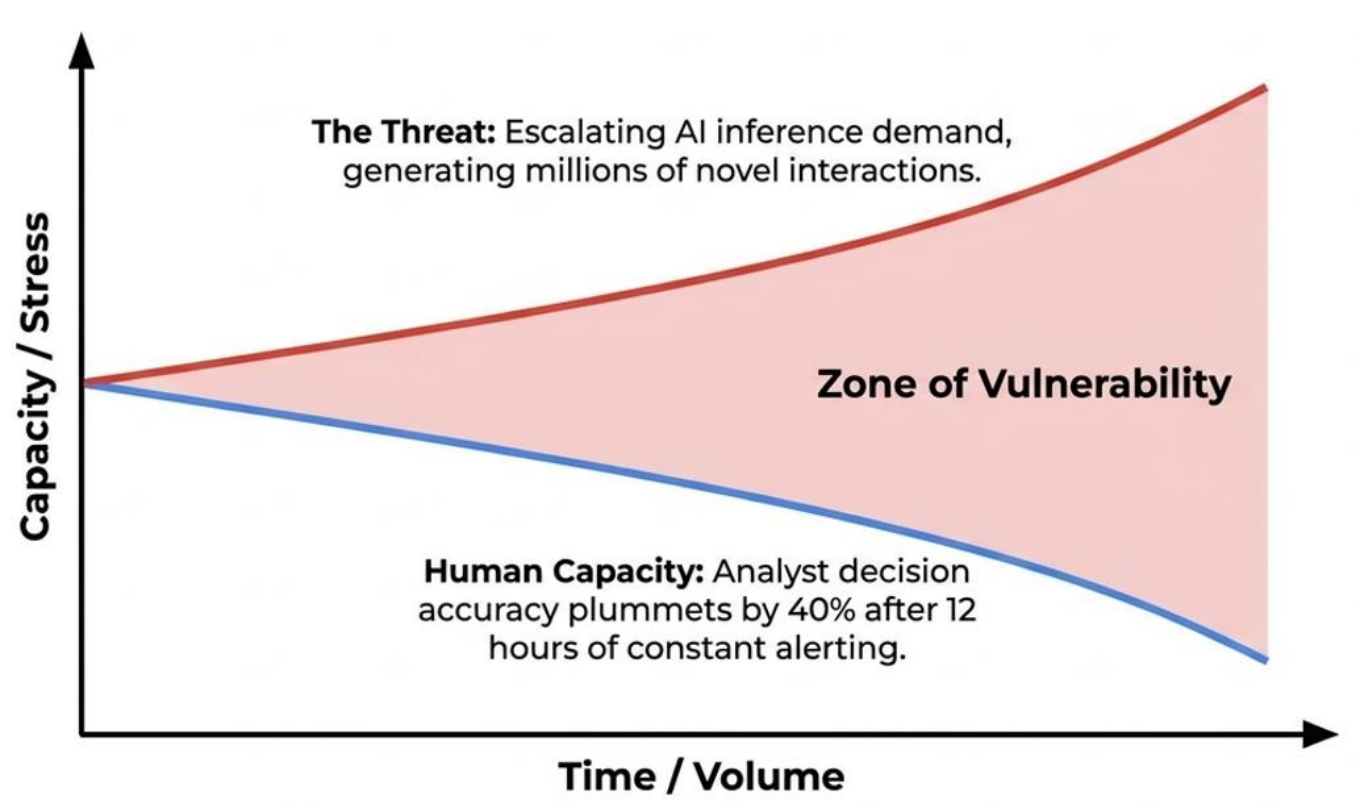

Defenders remain trapped in human-dependent processes

Legacy Architectures Cannot Scale to Machine-Speed

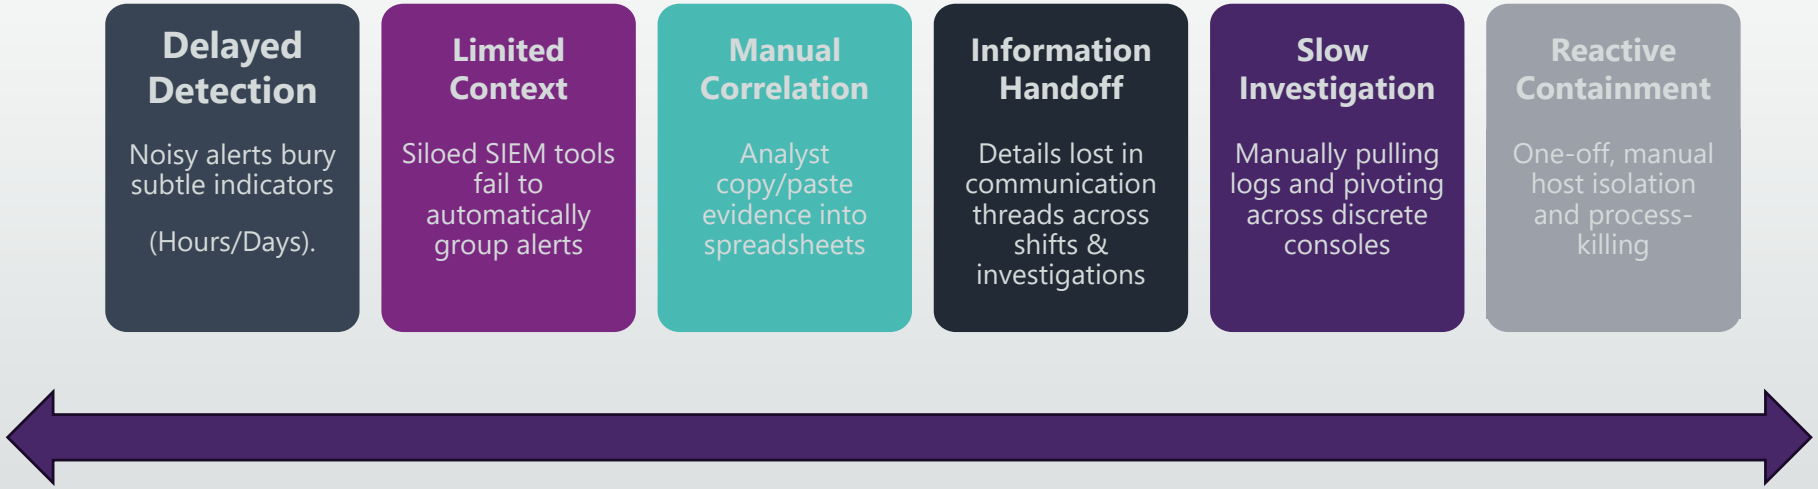

Bottom Line: Traditional incident response actions break when faced with novel and dynamic AI Threats

The Shift to Autonomous Operations

The Legacy Model

- Human-driven triage overwhelmed by high false positive rates
- Reactive, retrospective incident handling
- Siloed telemetry pipelines and fragmented investigations

The Unified AI Platform

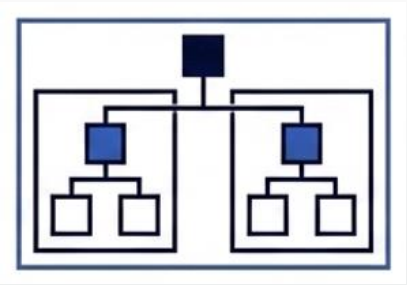

- Intelligent Data Foundation stitching telemetry together
- Machine triaged data leveraging behavioral ML
- Automated inline playbooks decreasing MTTD & MTTR

Evolving Automation: Hybrid Wins

Traditional Automation	Agentic AI	Hybrid Agentic AI (Goldilocks Solution)
Strengths - Efficient for known threats, predictable, consistent	**Strengths** - Autonomous planning, self-improving, adapts to new threats	**Strengths** Combines automation's precision with AI's adaptability. Operates within well-defined parameters while applying proven decision-making frameworks for novel scenarios
Limitations - Breaks on novel attacks, limited adaptability, requires constant manual updates	**Limitations** - Risk of policy violations, unintended errors, requires validation of guardrails	

AI is a Force Multiplier, Not a Replacement

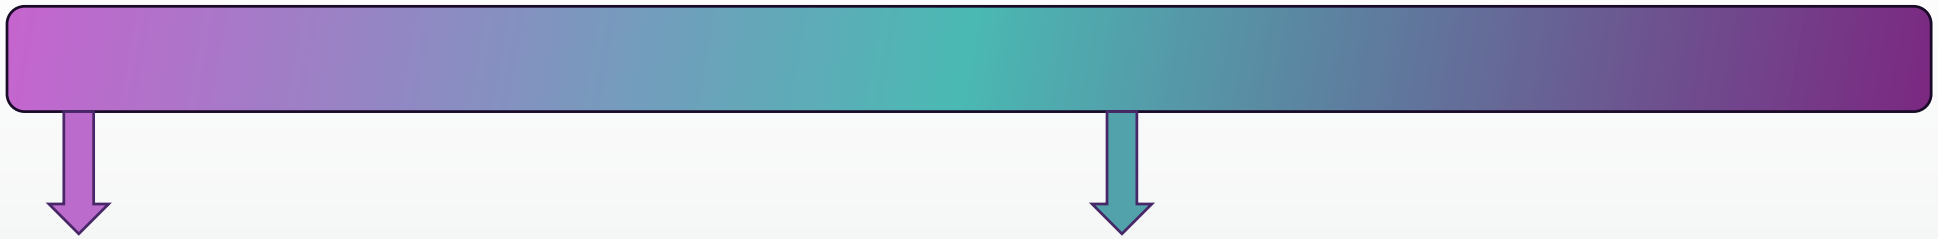

Where Machines Excel

- Rapid data processing and log aggregation
- Consistent execution of established playbooks
- 24/7 continuous monitoring without cognitive fatigue
- Pattern recognition across massive disjointed datasets

Where Humans Excel

- Navigating complex organizational dynamics
- Making nuanced risk trade-offs (e.g., business continuity vs. security)
- Creative problem solving
- Strategic threat hunting and security architecture design

Defender Perspective

- **Today actions start with the Analyst**
- **We need to flip the model**
- **Intelligence (AI), automation, and consolidation are critical for SOC's to be successful**

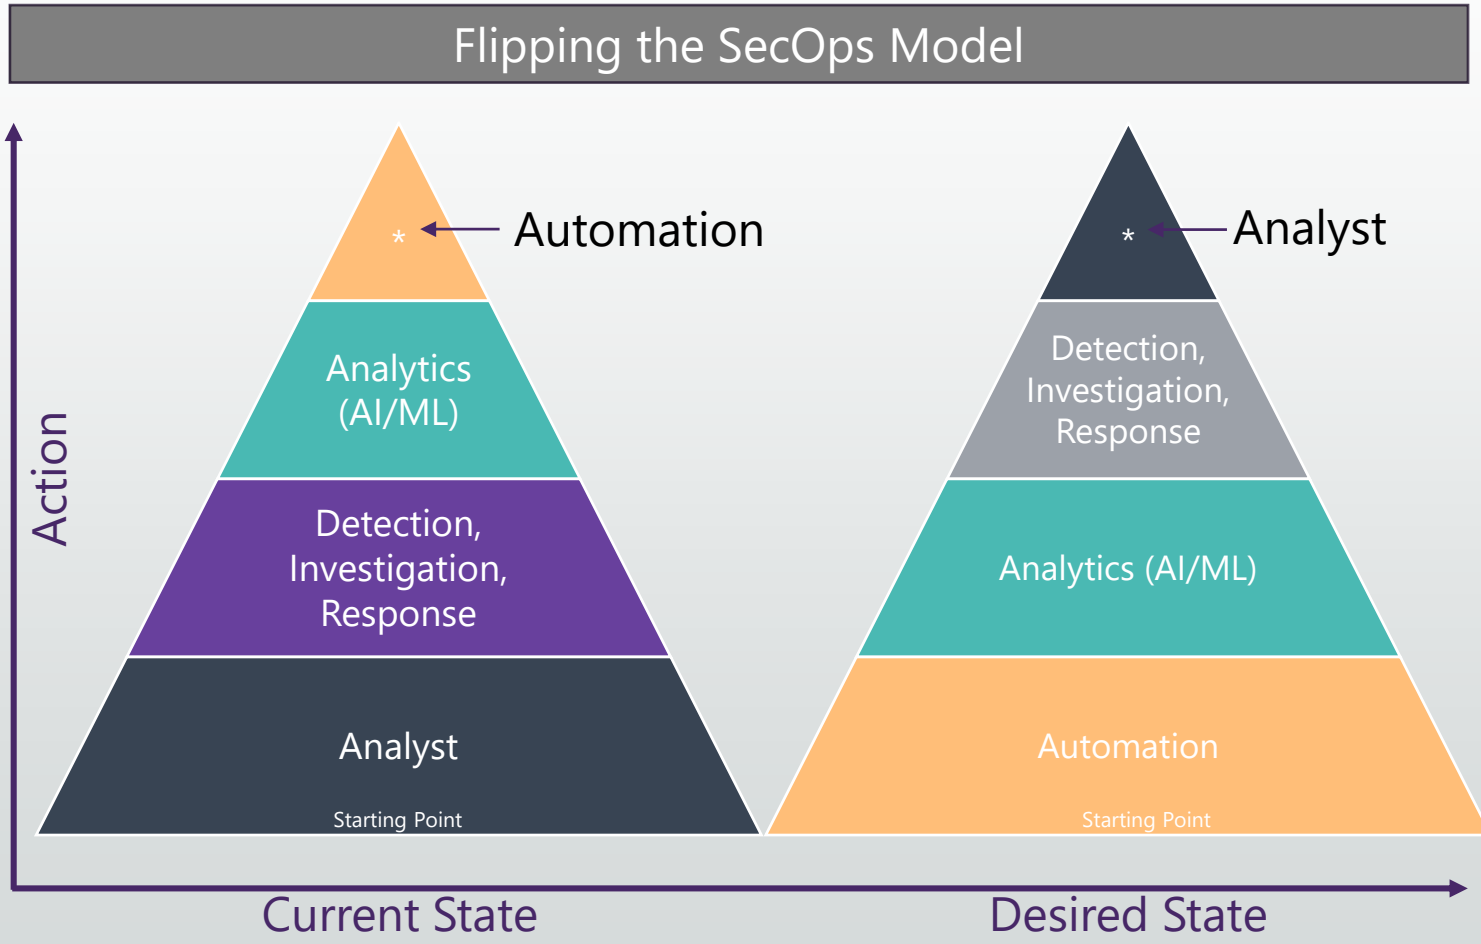

Elevating the Security Professional

Analyst -> Investigator

From:
Manually triaging thousands of basic alerts and fighting fatigue

To:
Investigating complex, novel threats and validating AI-generated root cause analysis

Engineer -> Architect

From:
Managing siloed telemetry pipelines and tuning noisy architecture alert rules

To:
Designing comprehensive methodologies and intelligent response workflows that AI scales organizationally

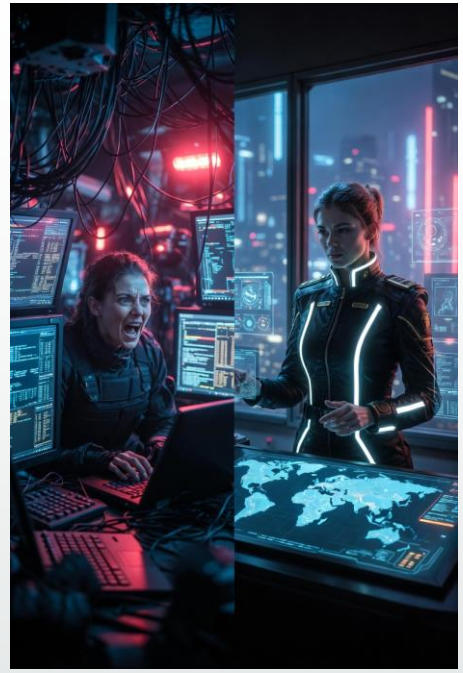

Responder -> Strategist

From:
Manually collecting evidence and executing one-off script containments

To:
Coordinating complex investigations, managing stakeholder communications, and proactive exposure management

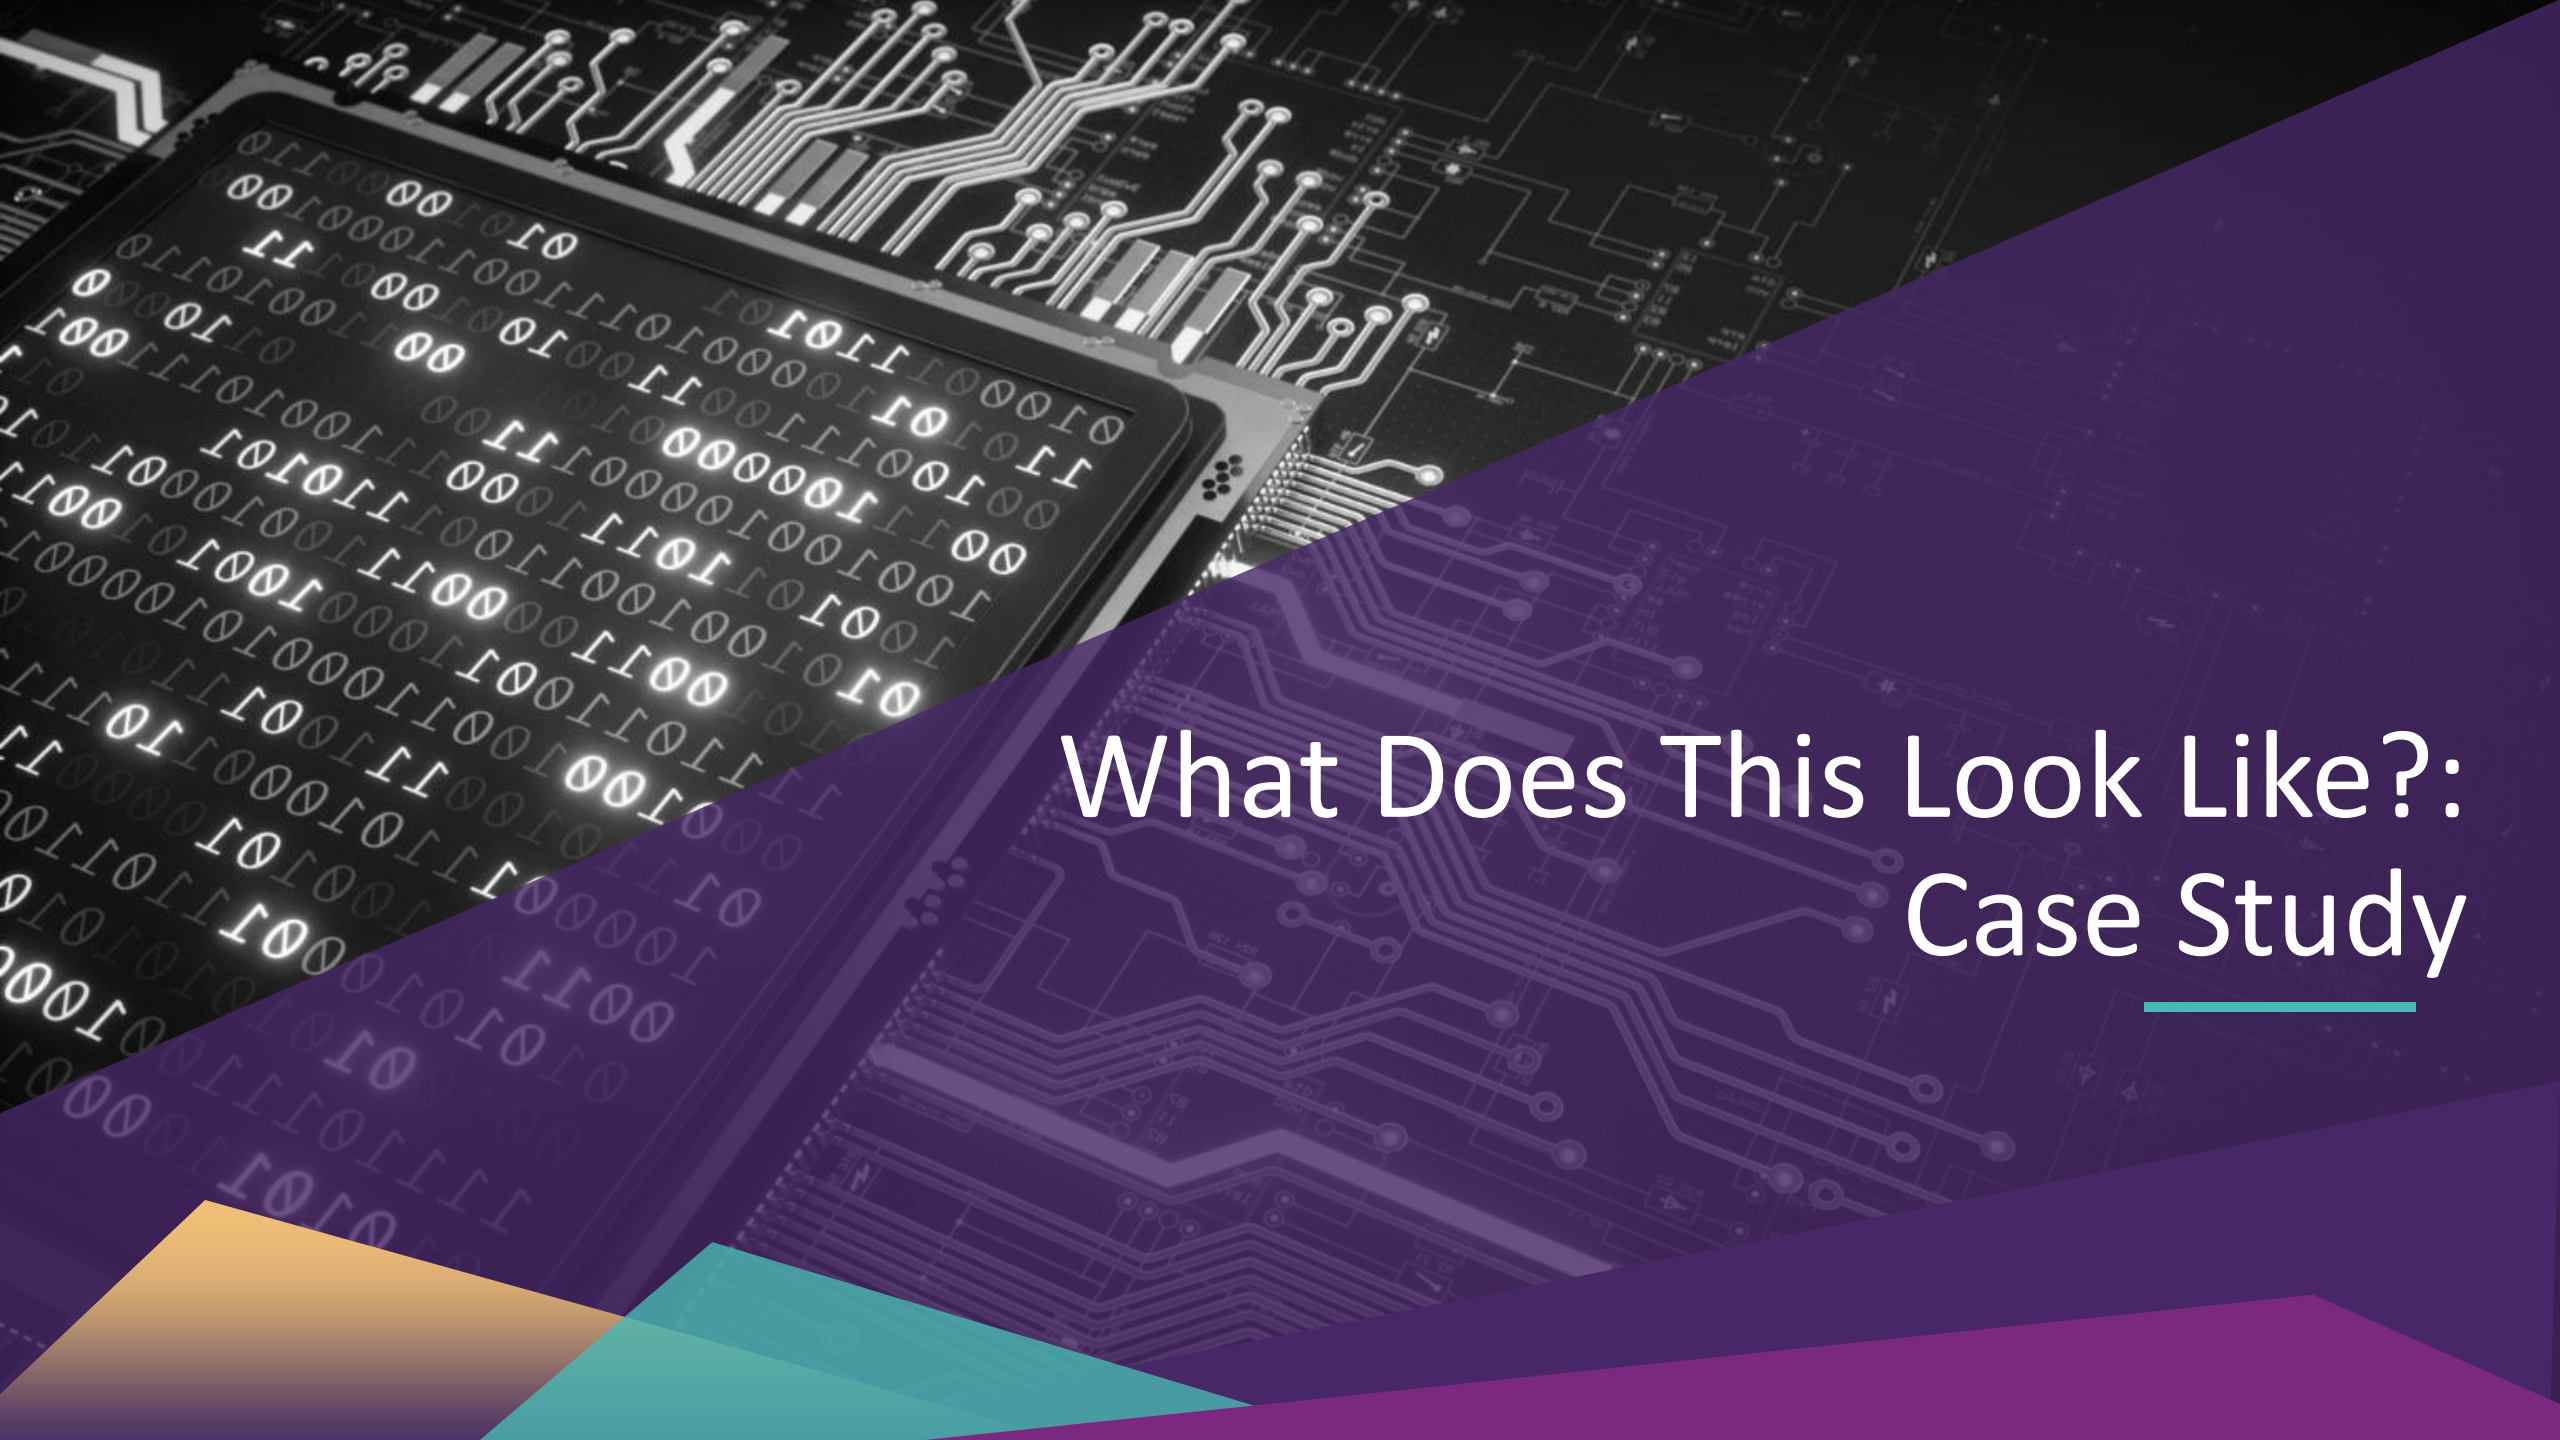

What Does This Look Like?: Case Study

The Goal

Dynamic Threat Intel Platform

- Run Threat Hunts On Demand
- Turn LLM output into usable, structured findings
- Keep execution local-first and private
- Avoid shared-session behavior and allow for additional use cases
- Output to operator-usable (aka me) dashboard

Stale Reports to Live Actionable Data

Multi-Source Threat Intel Roll-up Brief (v0)

Generated: 2026-03-03T23:43:15.422483+00:00

Data Sources Used

- KEV scored/correlated: `/Users/hammer/.openclaw/workspace/SecOps-Threat-Platform/data/normalized/2026-03-03/kev_scored_correlated_234315Z.json`
- CISA advisories normalized: `/Users/hammer/.openclaw/workspace/SecOps-Threat-Platform/data/normalized/2026-03-03/cisa_advisories_events_234312Z.json`
- NVD normalized: `/Users/hammer/.openclaw/workspace/SecOps-Threat-Platform/data/normalized/2026-03-03/nvd_events_234313Z.json`
- Unit 42 Tier 2 normalized: `/Users/hammer/.openclaw/workspace/SecOps-Threat-Platform/data/normalized/2026-03-03/paloalto_unit42_events_234315Z.json`
- CrowdStrike Tier 2 normalized: `/Users/hammer/.openclaw/workspace/SecOps-Threat-Platform/data/normalized/2026-03-03/crowdstrike_blog_events_234315Z.json`
- Cisco Talos Tier 2 normalized: `/Users/hammer/.openclaw/workspace/SecOps-Threat-Platform/data/normalized/2026-03-03/cisco_talos_events_234315Z.json`

Executive Summary

- KEV tracked: **1531**
- Corroborated KEV findings: **7**
- Corroborated preferred-vendor findings: **2**
- CISA advisories (high-signal): **5 / 30**
- NVD context events: **200**
- Tier 2 research events: **40** (high-signal: **6**)
- Mode: conservative, correlation-gated escalation

Top 10 Action Queue (Correlation-Weighted)

1. **CVE-2026-22719** | P1 | risk=100 | rollup=110.0 | corroborated=yes | vendor=Broadcom
2. **CVE-2023-46805** | P1 | risk=100 | rollup=100.0 | corroborated=no | vendor=Ivanti
3. **CVE-2025-49113** | P1 | risk=90 | rollup=100.0 | corroborated=yes | vendor=Roundcube
4. **CVE-2018-8639** | P1 | risk=95 | rollup=98.0 | corroborated=no | vendor=Microsoft
5. **CVE-2024-38094** | P1 | risk=95 | rollup=98.0 | corroborated=no | vendor=Microsoft
6. **CVE-2023-20269** | P1 | risk=95 | rollup=98.0 | corroborated=no | vendor=Cisco
7. **CVE-2023-36884** | P1 | risk=95 | rollup=98.0 | corroborated=no | vendor=Microsoft
8. **CVE-2022-41080** | P1 | risk=95 | rollup=98.0 | corroborated=no | vendor=Microsoft
9. **CVE-2020-3433** | P1 | risk=95 | rollup=98.0 | corroborated=no | vendor=Cisco
10. **CVE-2022-41082** | P1 | risk=95 | rollup=98.0 | corroborated=no | vendor=Microsoft

Corroborated Findings

- **CVE-2026-22719** | priority=P1 | sources=cisa_advisories, cisa_kev | boost=corroborated

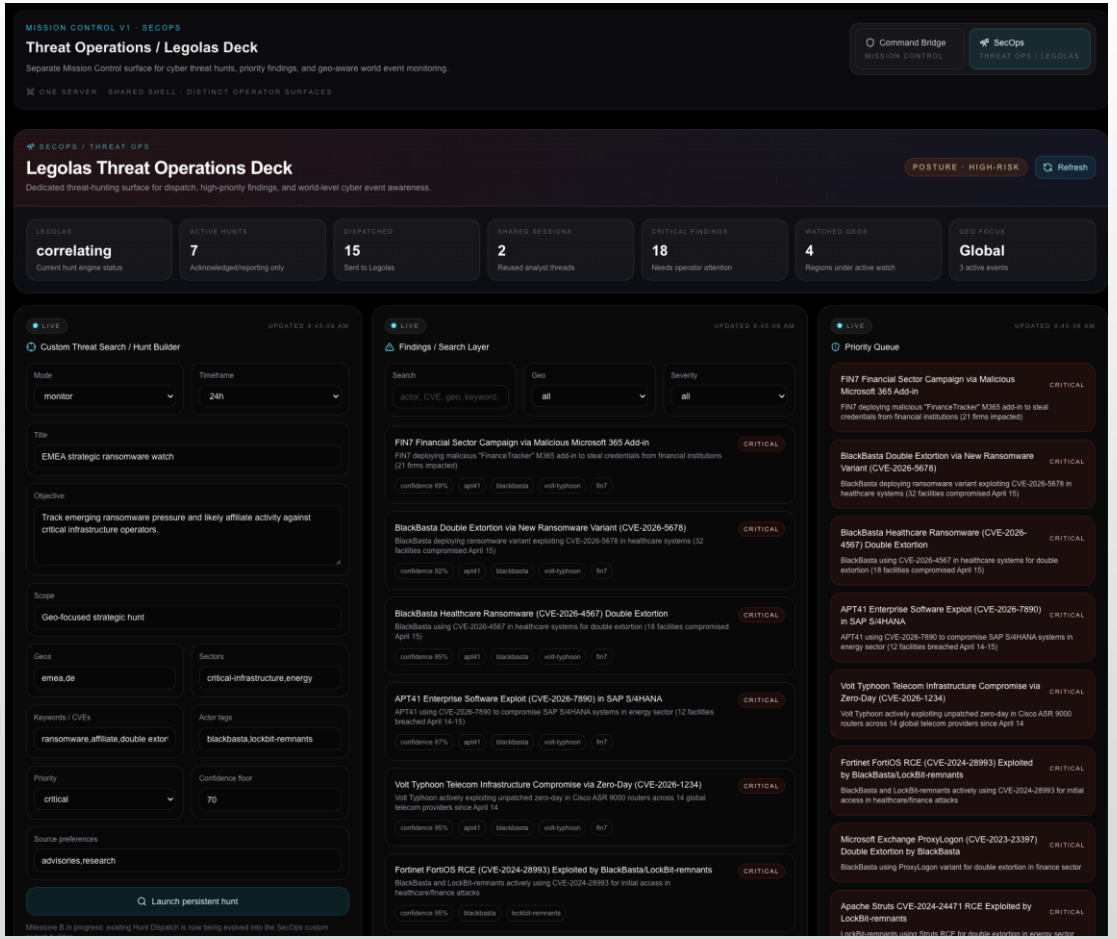

DEMO

Super Easy…

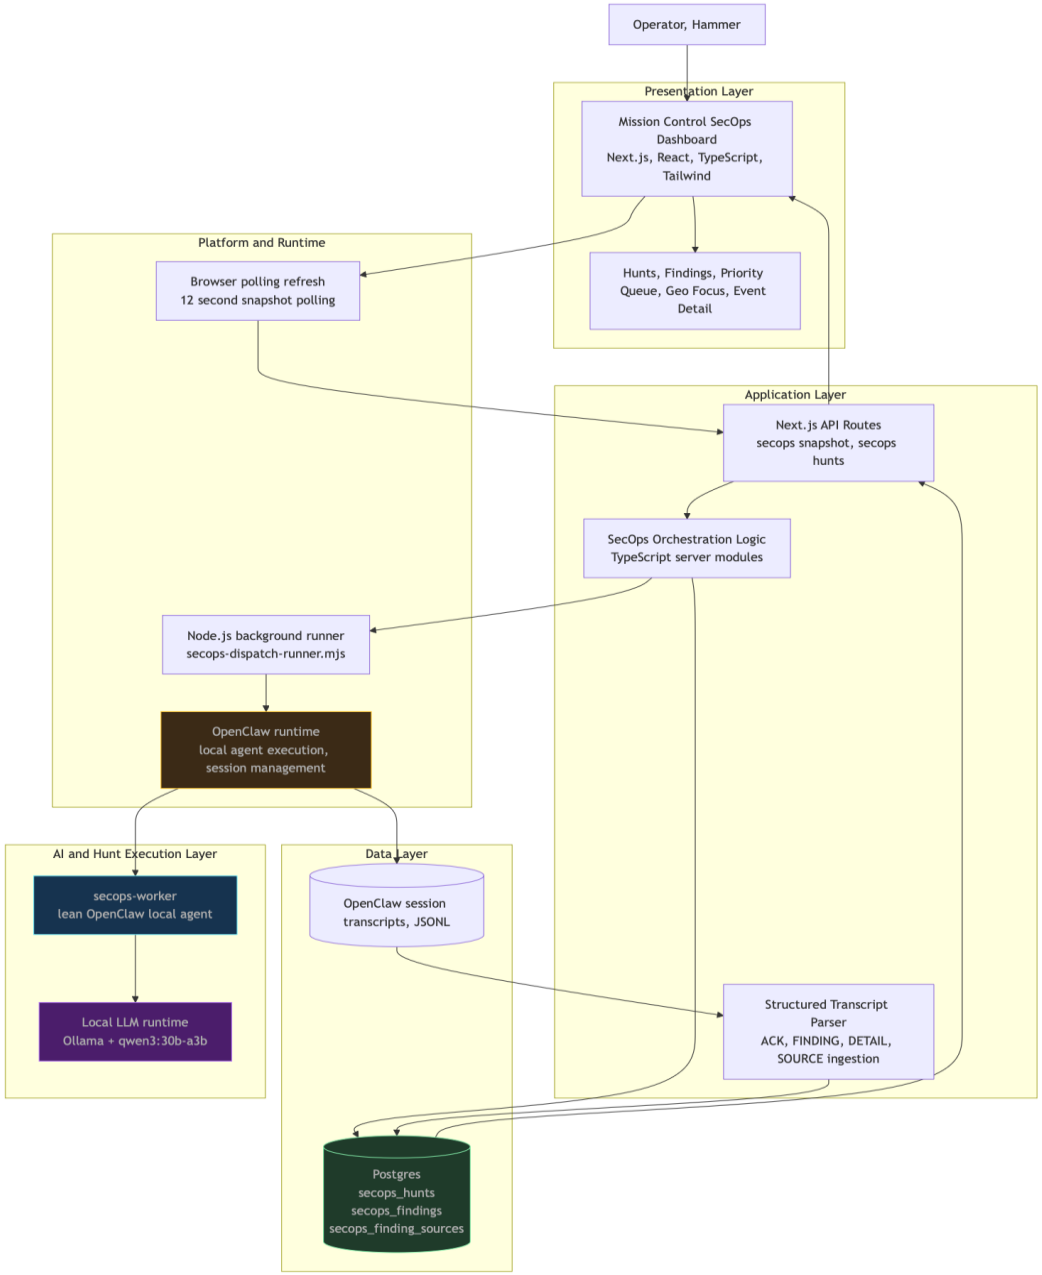

Isolation

docker

Code Repository

GitHub

LLM'S & Agents

OpenClaw

Ollama

ANTHROP\C

OpenAI

Presentation

Tailwind CSS

React

NEXT.JS

Orchestration

node.js

NEXT.JS

PostgreSQL

Data Layer

What Did I Learn?

5 Out of 100 Lessons…

- Start with a clear objective
- It's an iterative process
- Use the right model for the right task
- Things can get complicated (and expensive!) quickly…
- Have fun & learn!
- Ok.. One more… should you build or buy?

IT = Revenue Support

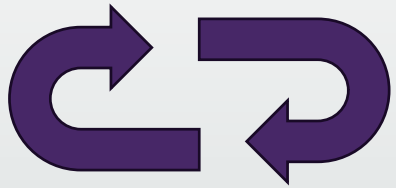

Security = Revenue Protection

You must be conscious of time and dollar investment

Final Thoughts

The Condensed Version…

1.) AI is going to continue to evolve and improve. This leads to benefits for Attackers AND Defenders.

2.) AI helps with but does not solve "Tool Sprawl". Look for opportunities to consolidate and focus on specific use cases.

3.) You don't need AI when scripting can solve the problem. Make smart SecOps decisions.

4.) Be aware of creating "AI Dependance". Humans still need to be in the loop.

Let us help.

AI Readiness Advisory – Overall look into AI and how it can impact your environment and preferred practices for rolling it out to your organization

Governance and Controls Assessment – Review of the governance program and controls in place and help implement AI guardrails to reduce risk

Tools Rationalization – Align tool capabilities to risk and control gaps. Identify where AI SecOps could broaden your organizations security posture

Contact your Account Manager

Thank you!

Questions?