ARAMBH LABS

CASE STUDY

Big Four professional services · Global MSSP

From 45-minute investigations to 5.3 minutes, across 20+ MSSP customers.

The managed security arm of a Big Four professional services firm runs SOCs for enterprises on a best-in-class stack, and still watched analysts spend most of an hour on every alert. Armor OS investigation agents now work the queue across 20+ of its customers, closing alerts in minutes with the evidence attached.

5.3 min
Mean time to investigate
Down from 45–60 minutes

20+
Customers onboarded
Onto agent-run investigations

100K+
Alerts in scope
Across the client base

78%
Closed as false positive
With the evidence attached

THE CHALLENGE

An hour per alert, at MSSP scale

As an MSSP, the firm runs security operations for many enterprises at once, each with its own environment, log sources, and escalation rules. Tooling was never the problem: client stacks run Google SecOps, Microsoft Sentinel, and CrowdStrike EDR.

The hour was the problem. For every alert worth investigating, an analyst still had to write the queries, pull the logs from each source, correlate them by hand, and write up a verdict. That is 45 to 60 minutes per alert, multiplied across a six-figure alert volume.

- 45–60 minutes of analyst time per investigated alert
- More than 100,000 alerts across the client base
- Every client environment different, so queries are rewritten per stack
- Top-tier tooling still left the correlation work manual

WHY ARMOR OS

A plan per alert. The agents build the investigation steps from the alert type and the specific customer environment it fired in.

Runtime execution. Queries written, logs pulled from each source, findings correlated live, not replayed from a pre-wired playbook.

Verdicts with evidence. Every alert comes back true positive, false positive, or suspicious, with the reasoning attached.

AT A GLANCE

INDUSTRY
Professional services (Big Four)

OPERATION
Global MSSP, 20+ customers onboarded

ALERT VOLUME
100,000+ across clients

CLIENT STACKS
Google SecOps, Microsoft Sentinel, CrowdStrike EDR

AGENTS
Armor Investigate

VERDICT MIX
78% FP · 12% susp · 10% TP

THE APPROACH

How the agents run an investigation

STEP 1

Plan

For each alert, the agents draft the investigation steps from the alert type and the customer environment it fired in. No two clients share a plan.

STEP 2

Execute

The steps run at runtime: the agents write the queries, pull the logs from Google SecOps, Sentinel, and CrowdStrike, and correlate what comes back.

STEP 3

Verdict

The alert closes as a true positive, false positive, or suspicious, with the correlated evidence attached for the analyst on record.

THE RESULTS

Across 20+ onboarded customers

Metric	Before	After
Mean time to investigate	45–60 min	**5.3 min**
Queries, log pulls, correlation	Manual, per alert	**Run by agents at runtime**
Alerts needing analyst judgment	Every investigated alert	**12%, flagged suspicious**
Verdict mix across the queue	—	**78% FP · 12% susp · 10% TP**

> *"The intelligent prioritization is game-changing. We went from drowning in 10,000 alerts to focusing on the 50 that actually matter.*
>
> Director, Large MSSP

See the same numbers on your own alerts.

Two weeks of your real queue, verdicts compared against your own analysts. Evidence attached.

Book a demo →