

CASE STUDY

From Blind Spots to Battle-Ready Detection

How a large enterprise closed its most dangerous detection gaps, cut security data waste, and cleaned up years of access sprawl — without buying a single new tool.

Client identity, dates, and figures anonymized for public sharing. Case composite reflects real engagement outcomes.

The Situation

A large enterprise running a modern endpoint detection and response (EDR) platform across **over 20,000 endpoints** believed its security posture was in good shape. Antivirus coverage was near-universal, the platform was well maintained, and the team had invested in a leading detection stack.


Leadership had one open question that no dashboard could answer: *could this environment actually catch a modern, largely automated attack* — the kind that moves at machine speed, uses legitimate tools, and doesn't drop traditional malware at all?

Counterveil was brought in to answer that question with evidence, not assumptions — and to fix whatever it found.

“We don't guess at maturity. We pull the live configuration, measure it against how modern attacks actually behave, and show you exactly where the gap is.”

What We Found

A full read-only interrogation of the platform's live configuration — prevention policies, detection telemetry, custom detection content, alert volume, and access governance — surfaced a pattern seen across many mature-looking environments: strong defense against yesterday's threats, and a wide-open gap against today's.



The behavioral controls that catch fileless, script-driven, and machine-speed attacks — the exact techniques modern automated intrusions rely on — were switched off fleetwide, even though the licensing to run them was already in place. Antivirus alone was carrying the entire defense.

At the same time, the alert pipeline was so overwhelmed with low-value noise that even a genuine detection would likely have gone unnoticed — and any AI-assisted triage layered on top would have spent its budget triaging noise instead of threats.

A parallel review of platform access found the identity story was just as overdue for attention: **nearly half** of all console accounts carried administrative privileges, a meaningful share hadn't been used in months (some never at all), and a small but high-consequence handful combined both — standing admin access nobody had touched in over a year.

How We Fixed It

Every recommendation was built around one constraint: this environment could not tolerate unplanned downtime, and no fix could be allowed to break a legitimate business process on its way to closing a security gap.

1. Diagnosed the real root cause first

- ✓ The core gap wasn't a missing tool or a budget shortfall — it was a configuration gap. The capability to close it already existed in the licensing the client owned. **Zero new spend was required** for the highest-priority fix.

2. Fixed the data before automating anything

- ✓ Recommended eliminating alert noise *before* layering on any AI-driven triage — every noisy alert an automated system reviews is wasted spend. Detection hygiene and cost efficiency turned out to be the same fix.

3. Built a zero-surprise rollout plan

- ✓ Every new control shipped in monitor-only mode first, then promoted to active blocking in small, reversible rings — canary, pilot, broad, full — gated on measured false-positive rates, with an instant rollback anchor at every stage. No control was ever switched on blind.

4. Delivered detection content tailored to this environment

- ✓ Rather than a generic rule catalog, new detection logic was mapped directly to the attack techniques actually showing up in this environment's own telemetry — credential theft, living-off-the-land execution, backup and log tampering — each shipped monitor-first with documented exclusions for legitimate admin and automation activity.


5. Found savings without touching what mattered

- ✓ A parallel review of the security data pipeline found real duplication and waste — but also flagged which high-volume feeds looked like waste yet were actually mission-critical for investigation and compliance. The savings recommended were the ones that were actually safe to take.

6. Handed over an execution-ready access clean-up

- ✓ Every over-privileged, dormant, or unused account was scored and organized into a ready-to-execute worklist — prioritized by risk, gated on human sign-off, and designed so nothing gets disabled in bulk or by mistake.

The Outcome



- ✓ A fully mapped, evidence-based picture of detection maturity — what's strong, what's missing, and exactly why it matters against modern attack techniques.
- ✓ A staged remediation roadmap the client's own team could execute with confidence, with rollback built into every step.
- ✓ A clear, defensible answer to “should we spend more on AI-driven security tooling” — not yet, fix the data first — that immediately paid for the engagement in avoided spend.
- ✓ A security data pipeline reviewed for waste without putting a single compliance-relevant log at risk.
- ✓ An access-governance program the identity team could act on the same week, not after another quarter of planning.

“The best security investment is often the capability you already paid for and never turned on. We find that first.”

Why Counterveil

We're a cyber defense consultancy built for organizations that can't afford a security recommendation that breaks something else. Every engagement is evidence-first, safety-gated, and built to be executed by your own team — not dependent on ours.

Evidence over assumptions. Every finding is pulled from live, read-only telemetry — never guessed, never templated.

Fix the configuration before the budget. We look for capability you already own before we ever recommend new spend.

Change-safe by design. Every control we recommend ships monitor-first, promotes in rings, and rolls back instantly.

Built for handoff. Our deliverables are worklists and roadmaps your team can run with — not a report that sits in a drawer.

Want to know what your own environment would show?

Let's talk about a detection-maturity assessment for your team.

www.counterveil.com

This case study is a composite drawn from real Counterveil engagement outcomes. Client identity, dates, and precise figures have been anonymized or generalized; described results reflect actual findings and deliverables.