Every core SOC function, covered. Your analysts own the verdict.

Dropzone AI gives each core SOC function to an agent that owns it end to end. The agents work the loop continuously and bring your analysts what needs a person, with the evidence already assembled.

The handoffs are the first thing to go

The hunt you scoped last quarter is still scoped. The advisory that landed Tuesday is open in somebody's tab. Nobody dropped either one: intel, hunting, detection and investigation each produce something another function needs, and connecting them is nobody's dedicated job. So the work gets done, it does not accrue, and next quarter starts roughly where this one did.

What an attacker can attempt in an hour, without deep skill, has been climbing. Coverage capped by human hours does not keep pace with a curve like that. Hiring does not answer it, because the constraint is how the work moves rather than how many people are moving it.

Dropzone AI gives each of those functions to an agent that owns it. Threat intel opens the loop, hunts carry it, investigations close it, and each finding hardens what the next pass looks for. Your analysts direct that work and own what it concludes.

We didn't build this because your team needed to work harder. We built this because the operating model needed rebuilding, and had for years.

WHAT IT IS TODAY		WHAT IT BECOMES
Coverage is rationed by hours	→	**Coverage is continuous**
People execute every step	→	**AI agents execute, people decide**
Work arrives as alerts	→	**Work arrives as findings**
Knowledge is tribal	→	**Knowledge is encoded**
Tools accumulate and people integrate them	→	**The agents carry the handoffs**
Trust is granted step by step	→	**Trust is set as bounds and checked after**

Until now, the work always outnumbered the team. Dropzone AI changes the math.

One agent per role, working as a team

Each agent does one job the way your best analyst would, then hands its work to the next one: intel to hunting, hunting to investigation, investigation to detection.

AI SOC Analyst

Know what happened

Investigates alerts end to end across your full tool stack, 24/7, showing the full evidence trail behind each verdict.

AI Threat Hunter

Know what you're missing

Runs hypothesis-driven hunts across your environment, turning what used to take 40 hours of manual hunting into about one.

AI Threat Intel Analyst

Know what threats matter

Condenses new threat reporting, scores what's serious and huntable, and drafts a hunt pack for your team to run.

AI Detection Engineer

Know where your coverage is thin

Continuously refines detections and drafts new ones tuned to your environment, keeping coverage current as things change.

What it means by role

CISO	Walk into the board knowing the environment was covered, not hoping it was. You decide what the agents handle on their own and what always comes to a person, and you can walk an auditor through what ran and why.
VP / DIRECTOR, SECURITY OPERATIONS	Alerts get investigated, threats get hunted and advisories get answered, so your team's hours go to the calls that need them.
SOC MANAGER	Coverage stops being a function of who is on shift and how deep the queue got overnight. The people you are trying to keep get to do the job they were hired for.
SOC ANALYST	The agents carry the groundwork and bring you finished work with the reasoning attached, so your shift goes to deciding. Your guidance shapes how the agents investigate next time, so your judgment is what keeps threats at bay.

WHY DROPZONE

One package, the whole SOC

You've had to decide which function gets help this year and which one waits for budget. Hunting usually loses. The package doesn't put you in that position: investigation, hunting, threat intel and detection engineering arrive together, right-sized to your organization, as a working team rather than four products you integrate.

The guidance becomes the asset

When an analyst guides a verdict or says what is normal in this environment, that judgment persists and shapes the next pass. Their expertise stops leaving when they do and starts compounding where the whole team benefits from it.

Enterprise scale, built at the regulated bar

Dropzone AI is built to the standard the most regulated teams set: **FedRAMP High (Class D) certified**, data residency, and full audit trails that hold up to government, finance and critical infrastructure scrutiny.

`SOC 2 Type II` `Single dedicated cloud tenant, no data co-mingling` `EU in-region deployment` `No customer data is used to train our models or our sub-processors' models`

The bounds are yours to set

You define what the agents may act on alone, where they stop, and what always comes to a person, and the audit trail shows those limits held. We meet the bar you'd hold any trusted partner to.

Every finding shows its work

Your analysts can follow the full chain of evidence behind a decision, review it, and guide it. The reasoning arrives with the finding, so reviewing it is a decision rather than a second investigation. Never a black-box score.

AI SOC Analyst

Know what happened

Alerts get a full investigation, and the shift goes to the calls that need a person

An alert takes evidence from four or five tools before anyone can call it. When the queue grows, teams decide what to investigate and what to close on instinct.

HOW AN INVESTIGATION RUNS

Alerts arrive from your SIEM, EDR, cloud, identity and email tools, with no data to migrate and no log normalization. AI SOC Analyst asks what happened, who it touched and whether it's normal here, then answers each question from your tools. It follows what it finds until it reaches a verdict. Your analyst gets that verdict with the reasoning in plain English, ready to review.

> **A finished case file you can open, guide and learn from.**
> The learning sticks. What your analyst teaches it shapes how the next investigation runs.

> **The same depth at 3 AM as on Monday morning.**
> Investigation quality stops varying with who is on and how deep the queue got overnight.

AI Threat Hunter

Know what you're missing

Threat hunting becomes a standing SOC function

Threat hunting runs on spare cycles and one analyst's initiative, so it's the first thing cut. The threats that never generated an alert stay unexamined, and nobody can say what was checked.

HOW A HUNT RUNS

Start from a hunt pack mapped to MITRE ATT&CK, or from your own hypothesis. AI Threat Hunter queries the SIEM, EDR, cloud and identity data that can answer the question, then works each lead to an outcome. What reaches your team is what was checked, what was found and why it matters, with the queries and evidence on the record.

> **Hunting keeps going on a busy week.**
> When the queue fills up, threat hunting isn't the first thing cut.

> **The findings aren't only attacks.**
> Hunt reports surface misconfigurations, shadow IT and vulnerabilities that were never going to generate an alert.

AI Threat Intel Analyst

Know what threats matter

The advisories that matter become hunts your team can run

More threat reporting comes in than anyone can read, and few advisories ever get checked against your environment. When leadership asks “are we affected?”, the answer depends on who had time to look.

HOW THREAT INTEL BECOMES A HUNT

New threat reporting from vendor research, government advisories and the security press is condensed into one story per threat and scored on how serious, exploitable and huntable it is. When a story is worth hunting, its hunt pack goes after how the attack works: the actor’s techniques and the vulnerabilities they exploit, with the known indicators as one hunt among several.

> **Your business sets the priority.**
> Stories are ranked by what matters to your business, so the loudest headline doesn’t automatically come first.

> **An answer with evidence behind it.**
> Every hunt comes back with what was checked and what was found, including visibility gaps and detection opportunities, so “are we affected?” has an answer you can show leadership.

AI Detection Engineer

Know where your coverage is thin

EARLY 2027

Your detections get better every week

Your detection set gets worked every week instead of whenever someone finds the time. Noisy rules get tuned rather than muted, rules that stopped earning their place get retired, and the recommendations come from what your own investigations found. You see the effect on alert volume before anything ships, and your engineers decide what does.

See the team work together, on your alerts.

dropzone.ai/request-a-demo