DATASHEET

NightBeacon CMD

An AI-driven SOC workbench, built inside a live SOC.

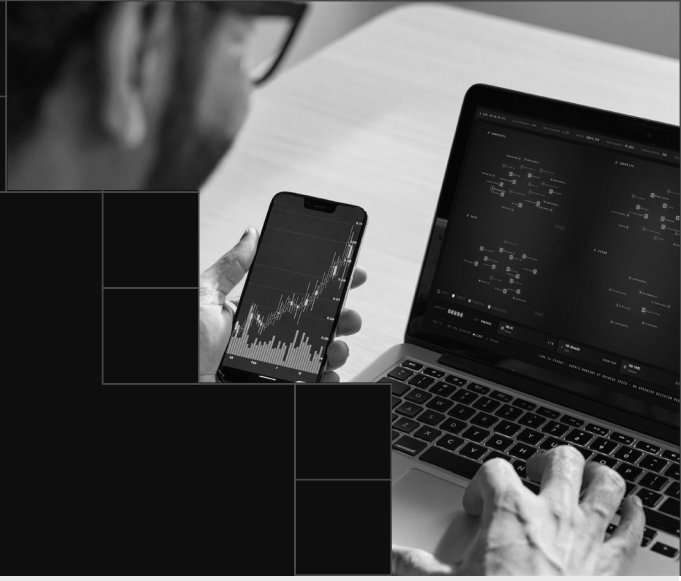

What It Is

NightBeacon CMD is an AI-driven SOC workbench: the same operating surface Binary Defense's own analysts use to run our live 24/7 SOC, now in your tenant for your team to run. It wasn't built in a lab. It was built and hardened inside a working SOC by the analysts defending real clients against real threats, so the workflow already fits the way a SOC actually operates.

It runs on a zero-queue model. Instead of handing analysts a bottomless list of alerts to triage and stitch together, NightBeacon CMD auto-correlates related alerts into a single situation before anyone opens a ticket. Your team works a short list of situations, not a flood of alerts, and each situation arrives enriched, correlated, and verdict-ready with roughly 80% of the investigation already done. This allows SOCs to spend the shift deciding and acting, not assembling evidence.

With triage handled at machine speed, the queue trends toward zero and your analysts get their time back for the work that moves the needle: hunting, tuning detections, and building playbooks. Every verdict is a glass box, a plain-language evidence chain rather than a black-box score. The platform learns from analyst decisions and never from your data, it's a an AI and a decision your team can stand behind in front of a board or an auditor.

~80%	116+	80+	8,700+
Of investigation pre-built, day one	Connectors across 9 categories, growing	Threat-intel sources enriched per indicator	Malware rules in the detection pipeline

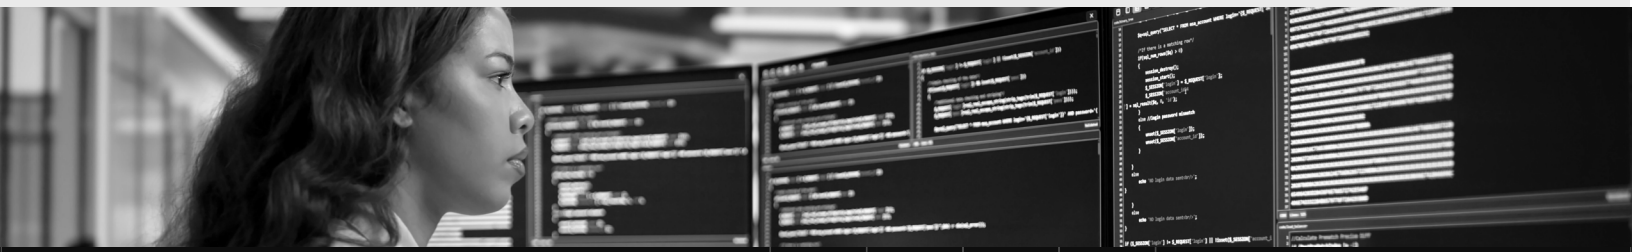

Not a lab experiment	Glass box, not black box	Human-in-the-loop by design
The real thing Built inside our live MDR on years of real analyst decisions across a large, diverse client base. Anyone can ship a model; no one else has the experience behind it.	Every verdict is explainable A plain-language evidence chain mapped to MITRE ATT&CK, down to which field the AI generated. Your analysts never have to say "the AI said so", it holds up in front of a board or an auditor.	AI analyzes; humans decide NightBeacon CMD does the analysis at machine speed; your people make every call. High-impact actions require analyst sign-off, a deliberate governance choice, and never trained on your data.

The Problem We Solve

Attackers move faster than any human team can keep up with alone.

Analysts are buried in alert noise and manual triage, so the real threats wait behind the false ones. Most analysts burn ~40 minutes on a single alert assembling context across a dozen disconnected tools before anyone can make a call. At 3–5k alerts per analyst, per day, and a false-positive rate that runs as high as ~99%, the real threats wait behind a wall of benign signal. You can't hire your way out of it. The only fix left is to shrink the work every alert demands.

Triage busywork Analysts pivot across tools for 20–40 min to assemble context before they can decide.	NightBeacon CMD pre-assembles the full, verdict-ready situation before the analyst opens the ticket.
Alert overload The same campaign hits the queue ten ways from five connectors; work scales with alert count.	Related alerts correlate into one situation, analysts work situations, not alert lists.
No time to hunt Teams are too buried in triage to hunt; when they do, it's manual and one platform at a time.	Hunt Assist runs a hypothesis, technique, or IOC sweep across every connected platform in parallel.
Black-box AI risk "The AI said so" doesn't survive a board meeting or an auditor.	Every verdict is explainable, a plain-language evidence chain mapped to MITRE ATT&CK.
Tool sprawl A dozen disconnected tools, each throwing its own alerts, no single operating surface.	116+ connectors across 9 categories, one workbench across the whole stack, no rip-and-replace.

The Value

Most tools give your analysts a longer alert list. NightBeacon CMD gives them a shorter.

The work-unit shift	The zero-queue model
Decide, don't dig Related alerts across every connector collapse into one situation, one verdict, one SLA, one set of response actions. The queue is bounded by real events, not alert volume.	Hunt, don't triage Alerts auto-correlate into situations before anyone opens a ticket, so there's no queue to grind through. With triage handled, your team goes hunting, tuning detections, and building playbooks.
The unfair advantage	**Calm by design**
The model, plus the signal Every pure-play vendor sells you a model. We sell you the model plus years of operational signal from a live MDR running real threats.	Silence you can trust Everywhere else, a blank queue means "what did we miss?" NightBeacon CMD's silence is backed by the numbers that account for it, quiet because the work is handled, not because it's hidden.

Client outcomes

What your team gets, each with the mechanism behind it.situation list.

~80%

Of investigation pre-built

Day-one connectors plus pre-assembled enrichment and correlation.

~46%

Faster situation summarization

NightBeaconAI auto-drafts the narrative before the analyst opens the situation.

24/7

Machine-speed coverage

The engine never rests, so your team can make decisions quickly.

- A situation list, not an alert list. Situations bound the queue by real events, not alert count.
- Fewer benign escalations. Per-customer false positive benign tuning and suppression rules surface only what's real.
- Explainable findings for boards & auditors. ATT&CK-mapped verdicts with a full, human-signed evidence chain.
- Catches novel, evasive attacks. Deep Scan detonation, behavioral analysis, and multi-layer signal fusion.
- No disruption from runaway containment. High-impact actions stay analyst-approved, with a controlled blast radius.
- Improves over time. Analyst verdicts feed continuous learning, with zero raw customer data.

Features & Benefits

Operations Room & Situations, your entire SOC on one screen.

If we show you one thing, we show you this: a cyber battle-management console, and the work-unit shift that makes a zero-queue SOC possible.

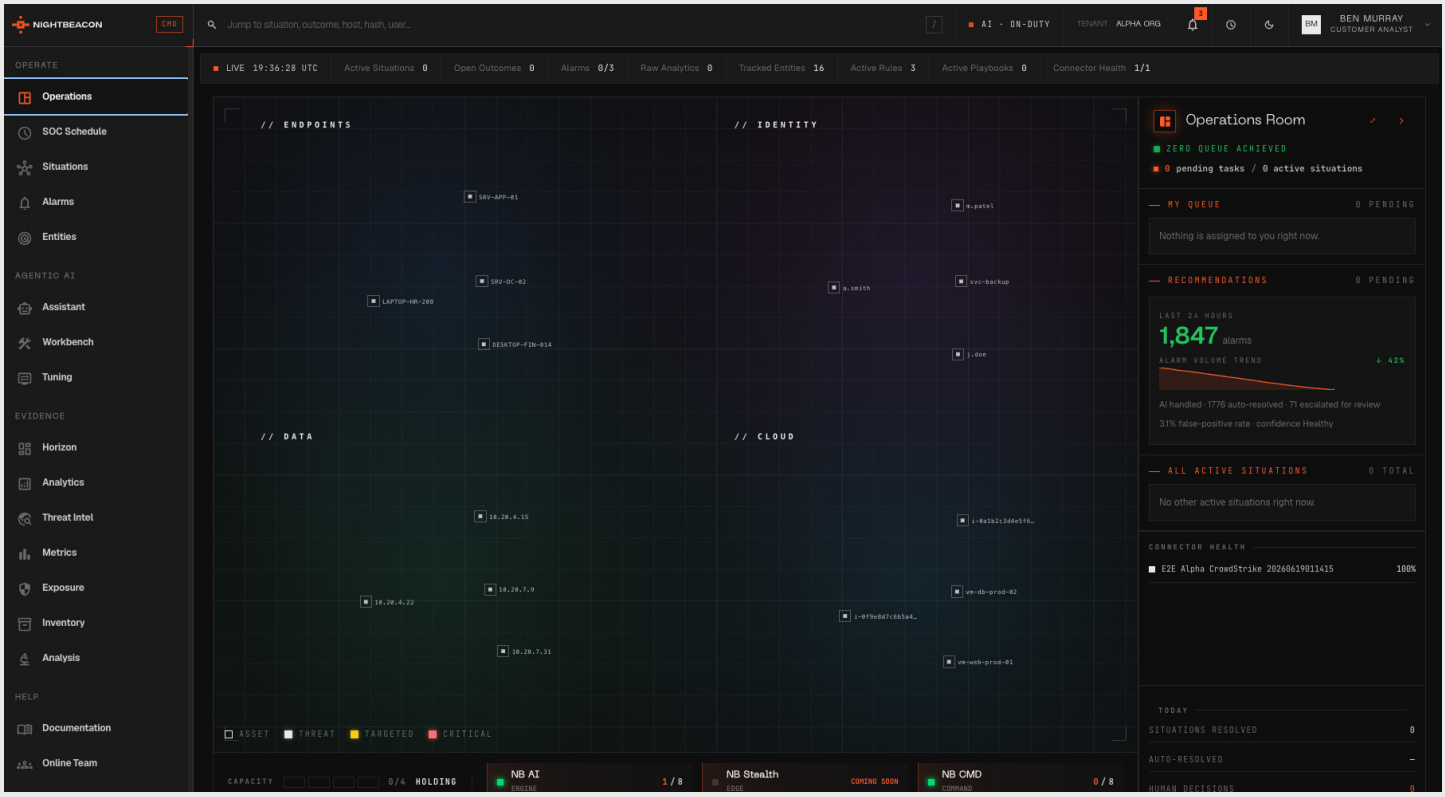

Operations Room

A live command console

A live entity map plots every asset across four quadrants, Endpoints · Identity · Data · Cloud, so an analyst clicks straight into the host, identity, or IP under attack. Alongside it runs the ranked situation queue, the AI agents on duty, and connector health, all updating in real time.

It's where every shift begins, the operations view an analyst opens the moment they log in. Click any situation to start investigating on the spot, launch it straight into a hunt, or see which analyst is on call, so the team is assessing and collaborating in real time, driving a faster time to response.

Asset Threat Targeted Escalated

Situations

The whole incident, as one story

Instead of thousands of alerts, an analyst sees a handful of situations. Open one and the AI has already correlated 50–60 disparate alerts across endpoint, identity, network, and cloud, by shared IP, user, or asset, into one narrated situation with a confidence score, a full timeline, and an interactive attack-chain graph.

From there the analyst follows a forward-only flow, the AI recommends the steps, proposes response actions, and surfaces tuning recommendations along the way.

Identify → Investigate → Verdict → Respond → Close

Silence as assurance

Peace-of-Mind Panel· quiet, quantified

When the queue clears, the sidebar flips from a work list to seven live signals, alarms in, auto-resolved, escalated, false-positive rate, confidence health, volume trend, and one-click tuning suggestions. The quiet isn't an absence of work; it's backed by the numbers that account for it. Everywhere else a blank queue means "what did we miss?", here it's quiet because the work is handled, not because it's hidden.

Benefit: Quiet becomes a status you can trust at a glance. The claim isn't "trust us, it's quiet", it's quiet, quantified, real and per-customer, every minute.

Cross-platform hunting

Hunt Assist

Type a hypothesis, pick a MITRE technique, or paste indicators. The agent translates intent into each platform's query language and runs the hunt across every connected platform in parallel for one consolidated answer with a consensus score, plan → act → verify → iterate.

Benefit: Proactive hunts across the whole stack at once that return signal, not noise, and won't melt the SIEM bill. It returns cited findings, and never claims a finding it can't back with evidence.

Sandbox & file analysis

Deep Scan & File Analyzers

Executables, Office docs (macro detection), PDFs, SVGs, and archives get disassembly, entropy checks, and C2-capability detection natively. Deep Scan detonates the unknown in an isolated sandbox, watches its behavior, and extracts the IOCs, including the C2 it beacons to.

Benefit: Answer "is this malicious?" on any file in-house, no reverse-engineering specialist, no extra tool licenses, plus IOCs to find every other infected host in minutes.

Agentic assistant · 5 modes + Auto

AI Assistant

One plain-English surface running on NightBeaconAI: Ask (security Q&A), Hunt (sweep all tools for a campaign), Analyze (verdict a URL, indicator, or payload), Act (propose a response for approval), and Build (generate a playbook), plus Auto to pick the mode.

Benefit: Triage, hunting, analysis, and response drafting in one place, grounded in your data and your verdicts. Ask "show me any alerts on this IP with evidence of ransomware", no one has to read it all.

Agentic investigation

Investigations Workbench

The agent and analyst share a desk. An LLM planner decomposes the situation into hypothesis steps, executes them where safe, builds a fully-sourced timeline, and proposes a verdict, before the analyst opens the situation. Unlike SOAR's fixed playbook, the agent writes its own plan per incident.

Benefit: Analysts confirm decisions instead of starting from scratch. When the attacker pivots, the investigation pivots with it.

Behavioral risk · UEBA

Entity Risk Radar

Every tracked host, user, and IP gets a behavioral risk score and an 8-axis spider graph, alert volume, severity, technique diversity, temporal pattern, observables, recency, network activity, and privilege usage, flagging exactly which axes are driving the risk.

Benefit: "Is this entity behaving abnormally?" answered at a glance, the UEBA picture without the manual queries.

Intel hub

Threat Intelligence

An aggregated surface for indicators (with confidence, source, and malware family), exploited CVEs flagged with KEV (CISA) and EPSS scores, threat actors, campaigns, and feeds, every indicator enriched against 80+ threat-intel sources, strictly passive.

Benefit: What's actually being exploited in the wild, in one dashboard, tied back to your alerts, with zero risk of tipping off an attacker.

Reports & dashboards

Report Center

One-click PDF / PowerPoint / DOCX packages, with NightBeaconAI writing each report in its audience's voice, Executive, Technical, or Risk & Finance, plus a business dashboard tracking hours reclaimed, dwell/MTTD, and cost avoided.

Benefit: The artifacts leadership opens and the dashboards analysts live in, board- and auditor-ready, the same data spoken in each reader's language.

Cross-org early warning

Campaigns & Supply Chain

Campaigns correlate the same attacker across three or more organizations in a rolling 72-hour window (refreshed every 15 minutes), surfacing affected-org count, shared MITRE techniques, indicators, and sightings, with cross-tenant privacy preserved through anonymized hashes. Supply-chain intelligence flags whether any of your own packages are implicated in an emerging compromise, like the recent wave of npm attacks.

Benefit: See a coordinated campaign or a poisoned dependency while it's still forming across the client base, not after it has already become your incident.

Built for regulated buyers

The governance model auditors, boards, and cyber insurers require.

Glass box	Your data stays yours	Human accountability
Every field, attributed AI-generated content carries the NightBeacon mark, so analysts see what came from raw data versus the AI, and can override it. The NightBeacon score sits beside the source severity; nothing is hidden.	Never trained on your logs NightBeaconAI learns from analyst verdicts via synthetic data. Your data exists only in memory during analysis, never persisted for training. zero customer data in the model, no cross-tenant exposure.	People own every decision NightBeaconAI does the analysis at machine speed; humans make every call. High-impact actions are approval-gated with a controlled blast radius, a deliberate design choice, not a limitation.

NightBeacon CMD puts the AI-driven SOC workbench Binary Defense's own analysts run every day inside your walls and under your team's control. Related alerts arrive already correlated into decision-ready situations, so your analysts stop assembling evidence and start hunting, and every verdict carries the reasoning to back it. Faster decisions, fewer missed threats, and AI you can always explain.

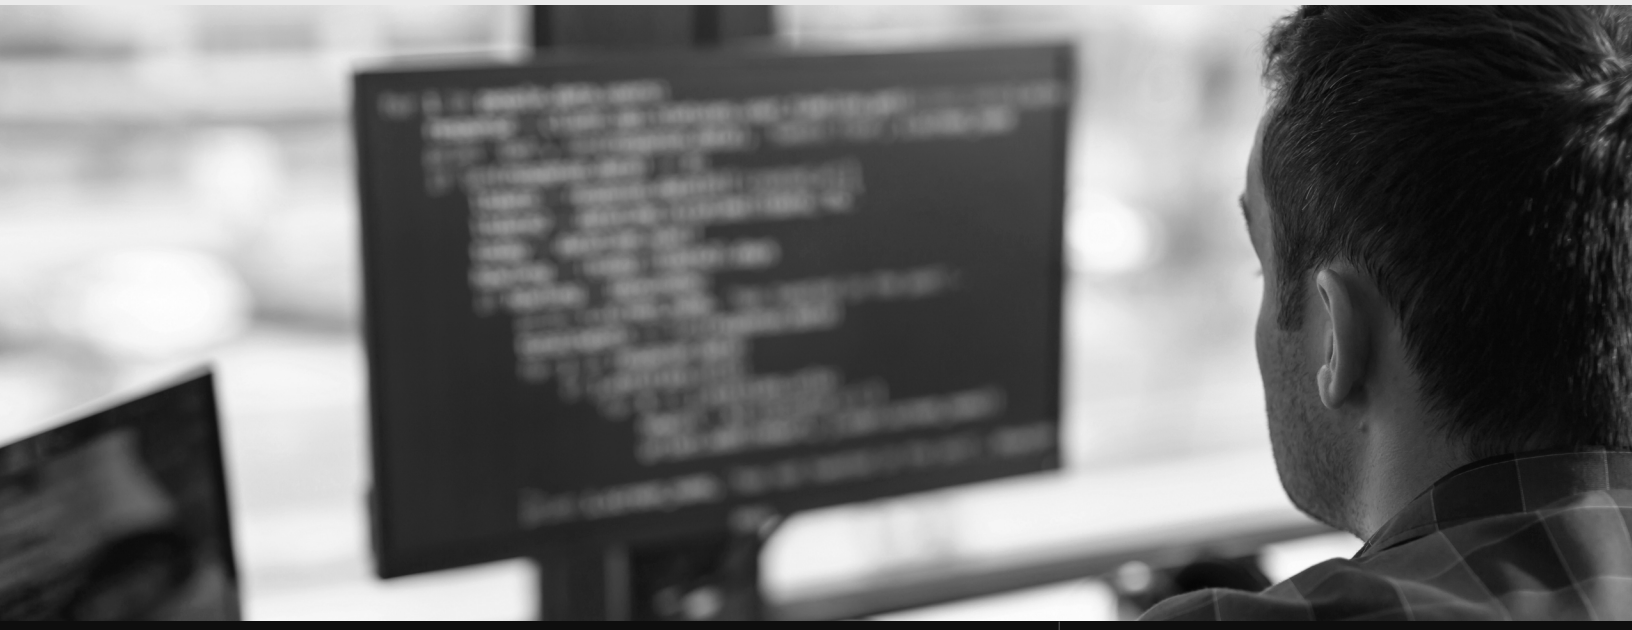

Binary Defense is a trusted leader in Managed Detection and Response (MDR), helping organizations across industries stay ahead of modern cyber threats. Our team of SOC analysts, threat hunters, detection engineers, and researchers works around the clock to deliver outcomes that matter: earlier threat detection, faster response, and stronger long-term security.

We take an attacker's mindset to defense, combining deep expertise with proactive tradecraft to protect what matters most.

Learn more at binarydefense.com, explore our blog for the latest insights, or follow us on LinkedIn.

600 Alpha Parkway, Stow OH 44224
sales@binarydefense.com

FORRESTER®
WAVE STRONG PERFORMER 2025
Managed Detection And Response Services

BINARY DEFENSE