CYBERSECURITY READINESS PODCAST SERIES

Episode 109

From Alert Fatigue to Agentic Defense: Redesigning the Human Role in the AI-Native SOC

with Joshua Weinick, AI Automation Engineer, Blink Ops

Host	Dr. Dave Chatterjee, Duke University
Guest	Joshua Weinick, AI Automation Engineer, Blink Ops
Topic Focus	Agentic AI in the Security Operations Center — Redesigning the Human Role in Automated Threat Detection and Response
Podcast Portal	https://www.cybersecurityreadinesspodcast.com/

Summary

In Episode 109 of the Cybersecurity Readiness Podcast Series, Dr. Dave Chatterjee is joined by Joshua Weinick, AI Automation Engineer at Blink Ops, to examine a question every security leader is now facing: what happens when the volume, velocity, and variety of cyber threats exceed what any human-staffed security operations center (SOC) can handle, and what role do humans play in a security function that increasingly depends on AI agents executing at machine speed.

Drawing on his background in cybersecurity consulting and his time embedded in Fortune 500 SOCs at Ernst & Young, Josh traces the limitations of legacy SOAR automation — rigid, Python-coded pipelines that broke whenever an edge case or a system change occurred — and explains why agentic AI, unlike its predecessors, is built to absorb change rather than collapse under it. The conversation moves through the three forces breaking the traditional SOC model — overwhelming alert volume, attack speed that has compressed response windows from hours to minutes, and the cognitive toll of analyst burnout — and arrives at a reframed model Josh calls "human after the loop," in which agentic systems take initial action within guardrails while humans review, audit, and reverse where needed.

The discussion examines what this shift means for emerging security talent, walks through a real-world 3 a.m. incident scenario that illustrates why a 45-minute human approval delay is now functionally equivalent to no response at all, and lays out five concrete steps organizations can take to begin the transition responsibly. Analyzed through Dr. Chatterjee's Commitment–Preparedness–Discipline (CPD) Framework, the episode reframes the SOC transformation not as a story about replacing security professionals, but about redesigning where human judgment adds the most value once the mundane, repetitive, and time-critical work is handled by governed AI systems.

Discussion Highlights

The Sustainability Crisis: Volume, Speed, and Cognitive Load

Dr. Chatterjee opens by naming the three forces that have made the traditional human-staffed SOC model unsustainable: a volume of alerts that has outpaced any team's capacity for meaningful review; an attack speed advantage that no longer allows time for a human in the loop to deliberate before damage occurs; and a cognitive toll — alert fatigue and analyst burnout — that is driving experienced practitioners out of the security profession altogether. He frames the question not as whether the

current model can be sustained, but as what should replace it, and what role remains for humans within it.

> *The most dangerous attacks today are designed to move faster than human decision cycles. The delay in getting a human to respond to an alert and take suitable action is just not acceptable, because by then the malware has probably penetrated the system and gotten itself deposited in multiple systems.*
>
> **— Dr. Dave Chatterjee**

From Human-in-the-Loop to Human-After-the-Loop

Josh traces his own path from cybersecurity consulting and Fortune 500 SOC work at Ernst & Young, where legacy SOAR automations were rigid, code-dependent, and prone to breaking entirely whenever an unanticipated edge case or a system change occurred — sometimes after seven or eight months of development work. That fragility, he explains, is what drew him to agentic AI platforms built to absorb change natively rather than require constant re-engineering. He introduces the concept of "human after the loop": rather than pausing an automated process to wait for human approval — a delay that can run 30 to 40 minutes when a manager is unavailable — agentic systems take action first, within defined guardrails, and a human reviews and can revert the steps taken afterward.

> *This isn't AI or agents coming to replace security people. It's that the place where humans are adding value to the security process is shifting. We're going from doing the work ourselves — logging into all the systems, copying and pasting, changing tabs — to actually designing how the work gets done.*
>
> **— Joshua Weinick, AI Automation Engineer, Blink Ops**

Turning Specialized Tools Into a Connected Stack

Josh describes how most security environments accumulate ten or fifteen point-solution tools, each generating its own detections without communicating with the others in real time. He explains how Blink Ops sits on top of that data layer, allowing tools such as endpoint detection and identity access management to share context automatically rather than requiring an analyst to manually cross-reference systems. In this model, the highest-value human role becomes architecting how the specialized tools connect and reason together, not operating any single tool by hand.

What the Shift Means for the Next Generation of Security Talent

Responding to a question about students entering SOC roles, Josh emphasizes that the most valuable skill is not technical tool proficiency but an understanding of the security and business context behind each action: why a given IP address is being enriched, what business continuity risk a given alert represents, and how the organization's different tools and data flows connect horizontally. He advises new analysts to understand how an existing ecosystem of tools can be orchestrated together rather than assuming new software must be built from scratch.

> *We're not just trying to close out as many alerts per day — we're trying to protect the business in a business continuity sense. Getting that fundamental, horizontal understanding of the security architecture is really, really important.*
>
> **— Joshua Weinick, AI Automation Engineer, Blink Ops**

The 3 A.M. Alert: When Minutes Replace Hours

Josh walks through a recurring real-world scenario: a global logistics firm's monitoring system flags a suspicious administrative login at 3 a.m. He uses the example to illustrate a structural problem rather than a tooling problem — that a 45-minute approval delay, once tolerable, now arrives well after an attack has already succeeded. Of the hundreds of daily alerts a typical SOC receives, often only one truly matters, and analysts exhausted by the other 99 frequently lack the clarity to catch it in time. Josh is direct that the failure is not the analyst's; it is a process built for an era when response time was measured in hours rather than minutes.

> *We've automated the easy part, which are the detections, and created these alarm systems that essentially deliver homework to exhausted analysts. The failure isn't on the analyst — the failure was the process itself. It's built for a world where you had hours, and now we're in a world where you only have minutes.*
>
> **— Joshua Weinick, AI Automation Engineer, Blink Ops**

The CPD Framework Applied to AI-Native Security Operations

Dr. Chatterjee applies his Commitment–Preparedness–Discipline (CPD) Framework directly to the SOC transformation discussed throughout the episode, mapping each pillar to a specific governance responsibility for organizations redesigning the human role in security operations.

CPD PILLAR	APPLICATION TO AI-NATIVE SECURITY OPERATIONS
COMMITMENT	Commitment begins with senior leadership explicitly recognizing that the human role in security operations must be redesigned — not defended in its current form. It means leadership actively engaging in deciding where humans belong in the loop, at the end of the loop, or designing the systems altogether, and acknowledging that detection-and-response speeds have outpaced what human decision cycles can deliver. This is a governance decision, not a technology procurement decision, and it requires leaders to mobilize organization-wide support and resources rather than delegate the question downward.
PREPAREDNESS	Preparedness means building the operational capability to strike the right balance between agentic AI execution and human judgment, and rehearsing that balance until it is proven effective. It includes defining clear, numeric thresholds for autonomous action, establishing guardrails around what an AI agent is permitted to do, piloting in recommended mode before granting autonomous authority, and continuously testing outcomes rather than assuming a well-designed system will remain effective without verification. Every part of the detection, triage, and remediation pipeline must be regularly reviewed and improved.

DISCIPLINE	Discipline is sustaining the model once it is built — keeping tools synchronized, keeping the human team trained to design and supervise AI systems rather than operate tools by hand, and continuously reviewing the data, guardrails, and reasoning that drive agent behavior. Organizations frequently build innovative processes that lose effectiveness over time because the motivation to revisit and reinvigorate them fades. Discipline requires treating the AI engine the way a car engine is maintained, except on a continuous rather than periodic basis, given how quickly attack speed and data quality can drift.

Five Steps to Begin the Transition

Asked how the CPD pillars translate into concrete action, Josh lays out a five-step roadmap: first, identify the highest-frequency, lowest-judgment work — password resets, onboarding and off-boarding, access reviews — where automation carries low risk and an obvious return; second, define explicit, numeric thresholds for what qualifies as high-confidence action an agent can take autonomously; third, establish guardrails defining the boundaries within which an agent is permitted to operate, often starting with read-only access; fourth, pilot the system in "recommended mode," where the agent proposes actions for human approval and earns expanded autonomy only as it demonstrates reliability; and fifth, invest in retraining the human team to design, supervise, and troubleshoot AI reasoning rather than execute tasks by hand. Josh stresses that the fifth step is not sequential but present throughout — automation that leaves the team behind creates a workforce chasing systems it does not understand.

> *If I'm talking to a board, I'd like to keep it very simple: AI handles the noise, so your experts can handle the judgment where it actually matters. None of this works if you automate and forget to bring the team along with you.*
>
> **— Joshua Weinick, AI Automation Engineer, Blink Ops**

Actionable Recommendations

RECOMMENDATION	DETAIL
Inventory High-Frequency, Low-Judgment Work	Identify the repetitive, monotonous tasks — password resets, onboarding and off-boarding, access reviews — that carry low risk and an obvious return on automation. This is the safest starting point for building organizational trust in agentic systems.
Define Explicit Autonomy Thresholds	Before automating any workflow, define in concrete, numeric terms what qualifies as a high-confidence action an agent can execute without human approval. Vague or implicit thresholds undermine both safety and leadership confidence in the system.
Establish and Document Guardrails	Define the boundaries within which an AI agent is permitted to operate — starting with read-only access and context-gathering before granting any ability to take action. Guardrails should expand only as evidence of reliable performance accumulates.

Pilot in Recommended Mode Before Granting Autonomy	Run new automated workflows in a mode where the agent proposes actions for human review rather than executing them outright. Let the evidence from this pilot period earn the agent's promotion to greater autonomy.
Shift From Human-in-the-Loop to Human-After-the-Loop for Time-Critical Alerts	For alerts where detection-to-impact windows are measured in minutes, redesign workflows so agentic systems can act within guardrails immediately, with human review and reversal authority applied afterward rather than as a precondition to action.
Retrain Analysts to Design and Supervise, Not Just Operate	Invest in retraining security analysts to architect, calibrate, and audit AI reasoning rather than manually execute tasks. When an agent makes a mistake, treat it as a missing data or boundary problem to correct — not evidence that the technology has failed.
Connect Point Solutions Into a Shared Reasoning Layer	Rather than replacing existing specialized tools, build a connective layer that allows them to share context automatically — for example, having an EDR platform check IAM permissions before escalating an alert — so no single tool operates in isolation.
Treat the AI Engine Like a System Requiring Continuous Maintenance	Unlike infrequent maintenance cycles for traditional systems, continuously review the data quality, reasoning, and outputs feeding AI-driven security decisions. Data quality is the oil in this engine, and it must be checked far more often than a periodic audit allows.

Time Stamps

TIME	CONTENT
0:02	Introduction — Dr. Chatterjee's background and credentials
0:49	Host framing: why the human-staffed SOC model is no longer sustainable
3:28	Joshua Weinick welcome and introduction
3:31	Guest career highlights — cybersecurity consulting, Ernst & Young, Fortune 500 SOCs, and the path to Blink Ops
6:36	The three forces breaking the SOC model: alert volume, attack speed, and analyst burnout
9:47	Josh's response: the shift from human-in-the-loop to human-after-the-loop
14:15	Why specialized security tools fail to communicate, and how a shared data layer solves it
15:34	Dr. Chatterjee reframes the conversation: replacing tasks, not people
18:06	Advice for students and new analysts entering AI-native SOC roles

20:58	Connecting the discussion to a holistic view of cybersecurity governance and readiness
22:00	The 3 a.m. scenario: a suspicious administrative login at a global logistics firm
24:51	The CPD Framework applied to AI-native security operations — Commitment, Preparedness, Discipline
30:29	Five concrete steps to begin the transition to agentic SOC operations
35:07	Dr. Chatterjee recaps the five steps and the engine/oil-change analogy for continuous AI review
38:24	Closing thoughts: the AI-native SOC, elevating the human role, and what separates future security leaders
39:55	Episode wrap-up and listener call to action

Memorable Quotes

> *What happens when the people we once depended on to fight cyber threats can no longer keep up with the volume, the velocity, and the variety of what is hitting them? The current model is not just failing at the edges — it is failing at the center.*
>
> **— Dr. Dave Chatterjee**

> *We've shifted from human in the loop, which is starting to fall apart from a solution standpoint, to human after the loop — where an agentic system is able to take action first, and then quickly someone is able to review what that agent had done and possibly revert one or more steps that it had taken.*
>
> **— Joshua Weinick, AI Automation Engineer, Blink Ops**

> *We are not talking about replacing people, we're talking about replacing tasks. We are examining the various security tasks that have to be carried out, and then making a call as to who will do it better — a human or a technology.*
>
> **— Dr. Dave Chatterjee**

> *The thing that'll separate leaders in the future is not going to be who has the most AI running around in their cybersecurity organizations. It's going to be who's able to rethink the human role the fastest — who moves their best people off of chasing alerts and onto governing and designing these systems.*
>
> **— Joshua Weinick, AI Automation Engineer, Blink Ops**

> *Managing the human infrastructure, the process infrastructure, the technology infrastructure in a sustained, consistent manner requires a certain discipline, and that is easier said than done. Companies start off with brilliant, innovative solutions, but over time, the motivation to review and reinvigorate those processes is no longer there.*
>
> **— Dr. Dave Chatterjee**

Dr. Dave Chatterjee | dchatte.com | dchatte@gmail.com | linkedin.com/in/dchatte