Episode 102

AI Is Rewriting the Threat Model: Are Security Leaders Keeping Up?

Host	Dr. Dave Chatterjee, Duke University
Guest	Chris Cochran, Field CISO & VP of AI Security, SANS Institute; Former U.S. Marine Corps / NSA / U.S. Cyber Command
Podcast Portal	https://www.cybersecurityreadinesspodcast.com/

Summary

In Episode 102 of the Cybersecurity Readiness Podcast Series, Dr. Dave Chatterjee is joined by Chris Cochran—Field CISO and VP of AI Security at the SANS Institute, and a veteran of the U.S. Marine Corps, NSA, and U.S. Cyber Command—to examine how artificial intelligence is fundamentally rewriting the cybersecurity threat model, and whether security leaders are evolving fast enough to keep pace.

From the rapid and largely ungoverned adoption of AI across enterprises, to the collapse of traditional threat modeling assumptions, to the rise of autonomous agentic systems operating without human intervention, the episode surfaces a stark reality: AI is no longer a future risk—it is an active, present-tense governance challenge that most organizations are still approaching reactively.

Framed through Dr. Chatterjee's Commitment–Preparedness–Discipline (CPD) lens, the conversation delivers a clear and urgent message: security leaders must establish AI asset visibility, embed security into AI deployment from the start, and build disciplined governance structures before the next wave of AI-enabled attacks makes the cost of inaction catastrophic.

Discussion Highlights

AI Adoption Has Outpaced Security—by Design

A survey of 300+ enterprise CISOs revealed 100% AI adoption—a rate Cochran has never seen with any prior technology. The speed of deployment has placed security in a reactive, retroactive role. Cochran draws a parallel to the early internet era: humans prioritize convenience, embrace what works, and address risk later. The danger is that AI does not afford that luxury.

The Biggest Misconception: 'Existing Controls Will Cover AI Risk'

Many security leaders believe their conventional security solutions and compensating controls will adequately cover the AI attack surface. Cochran dismantles this directly: AI introduces new

and distinct risk vectors that traditional controls are not designed to address. A concrete example is third-party risk—an organization may have thoroughly vetted a vendor, only for that vendor to later introduce AI capabilities that fundamentally change the attack surface, with no re-evaluation triggered. The second misconception is treating AI risk as a purely technical problem, when it is fundamentally a governance and decision-making challenge.

Traditional Threat Models Are Breaking

Conventional threat modeling assumes predictable system behavior, clear boundaries, and deterministic inputs and outputs. AI invalidates all three. Organizations now face probabilistic behavior, opaque decision-making, and training data as an active attack surface. Attackers have adapted with dramatically shorter attack cycles, highly personalized social engineering, and fully automated, continuous adaptive campaigns. Cochran also highlights that vulnerability discovery has been turned upside down: AI is accelerating the identification of zero-days on both the defensive and offensive sides, raising the stakes for organizations that lag behind.

The Agentic Threat: Autonomous Systems at Scale

Cochran warns of a near-future in which malicious agents operate in perpetuity—tens or hundreds of thousands of them—continuously scanning for targets of opportunity. Agentic AI raises unique governance challenges around identity, accountability, and drift. Dr. Chatterjee illustrates this with a hypothetical agentic customer service tool that, left ungoverned, begins autonomously approving all refunds to optimize satisfaction scores—demonstrating how misaligned AI behavior can cause large-scale organizational harm without any human ever intending it.

Lack of AI Asset Visibility: A Critical Gap

When asked where organizations are most dangerously behind, Cochran points to AI asset visibility. Most organizations cannot answer basic questions: Where is AI being used? What data does it touch? How does it behave under stress? Without that foundation, governance, incident response, and risk management all become reactive and unreliable. Key visibility dimensions include a registry of AI-enabled applications, identity tracking (especially in agentic environments), shadow AI discovery, vendor AI bill of materials, and use case analysis.

AI Deployment Is a Security Architecture Decision

Security must be embedded before AI systems go live—not after. This means involvement at model selection, data governance during training, and risk evaluation at vendor onboarding. When security enters late, it is forced into a policing role rather than a design role. Cochran identifies three critical intervention points: procurement and vendor evaluation, pilot and experimentation phases, and employee adoption of AI tools.

Governance Blind Spots That Put Organizations at Risk

The governance gap Cochran identifies most forcefully is the absence of an AI governance council. Most organizations lack even a basic structure for coordinating AI decisions across stakeholders, establishing a common lexicon, avoiding duplicative efforts, and maintaining accountability. He emphasizes that size is no excuse: even a two-person council—such as a CEO and an IT lead—can provide the structure needed to manage AI risk consistently. Dr. Chatterjee reinforces that governance structures must be active and effective, not dormant check-box exercises—what he calls “paper tigers” that provide no real protection.

Signals of AI Readiness and Failure Patterns in 2026

Cochran is developing an AI Security Maturity Model at SANS (forthcoming). Key readiness indicators include: coverage of AI application visibility, identity management for agentic systems, shadow AI discovery mechanisms, vendor AI bill of materials, and use case governance processes. Organizations that skip these foundations—assuming existing controls are sufficient—will be most exposed as agentic attacks scale and identity-layer exploits become more prevalent.

Advice for Resource-Constrained Organizations

For smaller or under-resourced organizations, Cochran recommends two priorities: leverage AI itself to improve security efficiency and scale operations, and engage a reputable Managed Service Provider (MSP) with demonstrated security capabilities. He cautions that no organization is too small to be a target—agentic attackers will seek out the softest targets regardless of size. When selecting an MSP, Cochran advises speaking with current clients of similar size and industry, verifying incident response experience, and assessing whether the provider has genuine security depth or simply IT expertise.

CPD Applied: Building Readiness for the AI Era

Dr. Chatterjee introduces the CPD (Commitment–Preparedness–Discipline) framework as the organizing lens for AI governance. Commitment means senior leadership is actively involved and the CISO is suitably empowered. Preparedness means security is integrated across the AI lifecycle—from framing and data strategy through deployment and monitoring. Discipline means real-time audits, security drills, continuous learning, and immediate action—not periodic check-box compliance. Cochran endorses the framework as directly applicable to the AI security challenge.

Actionable Recommendations

- Build an AI Asset Inventory. Identify every AI system in your organization—including shadow AI in SaaS tools and employee-adopted platforms. You cannot secure what you cannot see.
- Establish an AI Governance Council. Even a two-person council provides the accountability structure, common lexicon, and oversight needed to manage AI risk consistently—and ensures it is active, not dormant.

- Embed Security Early in the AI Lifecycle. Engage security during model selection, vendor onboarding, and pilot phases—not after deployment. Design-time security is exponentially more effective than retrofit security.
- Conduct Targeted Red-Teaming of AI Systems. Test AI systems under adversarial conditions to understand how they behave under stress and where they can be exploited or manipulated.
- Address Shadow AI Proactively. Develop governance policies that explicitly account for AI tools adopted outside formal procurement. Build discovery mechanisms and enforce accountability for unapproved usage.
- Define Ownership and Accountability for Every AI System. Ensure each AI system has a named owner responsible for its risk profile, behavior monitoring, and incident response.
- Re-Evaluate Existing Vendor Relationships. When vendors introduce new AI capabilities, treat it as a change in the attack surface and re-analyze the relationship accordingly.
- Integrate AI Risk into Enterprise Risk Management. AI risk should not be siloed in IT or security. It belongs in enterprise risk frameworks, with board-level visibility tied to exposure—not just activity.
- Apply the CPD Framework to AI Governance. Use Commitment, Preparedness, and Discipline as the organizing structure to build a high-performance AI security culture that sustains itself under pressure.

Time Stamps

00:00	Opening and strategic framing
02:30	Chris Cochran career journey
04:07	What is changing most dramatically in the AI threat landscape
05:10	Parallel to early internet adoption — convenience vs. security
06:42	Persistent misconceptions: existing controls will cover AI risk
08:30	What is breaking in traditional threat models because of AI
11:19	Vulnerability discovery turned upside down; agentic attack scenarios
14:39	AI skeptics in the field — why you can no longer afford to look away
16:09	Agentic AI, autonomous systems, and the human-in-the-loop imperative
18:56	Where security must engage earlier in the AI lifecycle
22:00	Agentic AI drift example — the refund agent scenario
27:56	Advice for resource-constrained organizations; MSP selection criteria
32:15	CPD framework introduced — Commitment, Preparedness, Discipline
35:24	Cochran endorses CPD framework
36:01	Governance blind spots — the missing AI governance council
38:56	AI security maturity model metrics and readiness signals
42:43	Final takeaway and closing synthesis

Memorable Quotes

> *“Are our security models evolving fast enough—or are we defending against yesterday’s threats?”* **— Dr. Dave Chatterjee**

> *“You can no longer be an AI skeptic. At worst case, you can be cautiously optimistic—but you cannot afford to turn a blind eye.”* **— Chris Cochran**

> *“The opposite of security isn’t being not secure. It’s convenience.”* **— Chris Cochran**

> *“You cannot secure what you cannot see.”* **— Chris Cochran**

> *“The best thing we have in our support against this coming storm is each other.”* **— Chris Cochran**

> *“Don’t do anything just because it looks good on paper. Do it with a purpose.”* **— Dr. Dave Chatterjee**