EPISODE 16 | Webinar Recap

Agentic AI for Fraud & AML

April 27, 2026

This transcript was auto-generated and may contain errors or inaccuracies.

Jeff Welcome to "What the F Happened? Fraud and Financial Crime Deconstructed," a DEFEND podcast where we break down what's actually happening across fraud scams, AML and financial crime.

Emily Each episode cuts through the noise to explain the tactics, trends, and real world impact behind the headlines so you're better prepared for what comes next. Let's get into it.

Jeff Right off the bat today. Um, I want to hit you with a number that should honestly just chill. Anyone paying attention to the financial sector?

Emily Oh, absolutely. It's a staggering figure.

Jeff Four point five times. Yeah. That's, uh, according to Interpol, criminals using autonomous AI agents are making traditional fraud four and a half times more profitable.

Emily Right.

Jeff So this asymmetric cyber war everyone keeps warning us about. I mean, it is already here.

Emily It is. And the adversary is, you know, they aren't waiting for compliance approvals. They're operating at an industrial scale right now, completely unburdened by the, uh, the regulations that financial institutions have to navigate.

Jeff Yeah. Which is terrifying. Yeah. Welcome to today's deep dive, everyone. We are unpacking why so many of these Agentic AI deployments are just crashing and burning, specifically in the, you know, hyper regulated world of fin crime fraud and AML.

Emily Right? Anti-Money laundering.

Jeff Exactly. And our mission today is to cut through all that vendor hype. We want to uncover why conversational AI agents are becoming the new gold standard for catching these guys. Mhm. Um, how they actually differ from your basic chatbots and why dropping a powerful AI into a highly regulated environment without the right guardrails is just a recipe for disaster.

Emily A total systemic disaster. Yeah. Yeah. To understand why these deployments fail, we have to look at the pressure cooker they're operating in. Okay. When we say criminals are making fraud four and a half times more profitable, we need to look at the actual mechanism of that attack.

Jeff Right. How are they doing it?

Emily Well, bad actors are using AI to autonomously scrape the dark web, stitch together leaked credentials, and generate synthetic identities by the tens of thousands. Wow. Yeah. They use these agents to simultaneously probe a bank's defenses, and they automatically recalibrate their approach the exact millisecond they hit a threshold block.

Jeff So it's basically like they're running high frequency algorithmic trading, but for financial crime.

Emily Exactly. That's a great way to put it.

Jeff Meanwhile, if you look at the traditional fraud and AML operating model from, what, five years ago, it's entirely broken. It relies on reactive responses, manual rule updates, investigations that sit in a Jira queue for like two weeks.

Emily Right. Waiting for an engineer.

Jeff Yeah. It's like, um, it's like showing up to a modern cyber war with a seventeenth century musket, right?

Emily It's totally outmatched.

Jeff I mean, for fraud teams listening, you know how uncomfortable this reality is. The manual tools simply do not match the speed and scale of modern threats.

Emily What's fascinating here is that the industry actually agrees on this disparity. Like the consensus overwhelmingly points toward agentic AI as the necessary countermeasure. Right.

Jeff Everyone's talking about it.

Emily But agreeing on a buzzword doesn't mean agreeing on the execution. A software vendor slaps the label agent on a basic script. Institutions buy it expecting a silver bullet, and then, uh, they find out it can't handle the complexity of actual banking environments.

Speaker 3 Yeah, we really need to.

Jeff Strictly define the terminology here about the differences between AI chat, AI agents, and conversational AI agents.

Emily It's a crucial distinction.

Jeff So AI chat is the baseline, right? Good for retrieving info or summarizing. You prompt it like, hey, summarize this suspicious activity for this account and it just spits out text or code, right?

Emily It informs decisions, but it cannot execute them. The humans still has to take that text and do something with it.

Jeff Exactly. I kind of picture AI chat as like the intern who hands you a really good research report. It's helpful, but you still have to do the work spot on.

Emily Now, moving up from that, you have AI agents. These are autonomous systems that actually take action. Okay? They move from insight to execution. So instead of just flagging an account, an AI agent might autonomously adjust risk thresholds or place a hold on funds based on its own analysis.

Jeff So that's like an aggressive intern who just goes and does stuff in the background without asking.

Emily Exactly, which, you know, is a bit scary.

Emily Yeah,

Jeff And briefly contrasts building this yourself versus buying it.

Emily Yes. Building your own takes forever, right? The slow time to value is rough, but using pre-built domain trained agents that already understand fin crime workflows, they get it immediately. Which brings us to the Holy Grail, Conversational AI agents. These combine the plain language prompts of the chat with the complex task execution of the agent.

Jeff So a highly skilled collaborator sitting right next to you.

Emily Exactly. You could ask a system to build a rule to detect high risk ACH credits tested on last month's data and deploy if performance improves.

Jeff Wow.

Emily Yeah, the system actually creates the rule, back tests it, evaluates it, and deploys it all from that one conversational prompt.

Jeff Okay, so if agents are so incredibly powerful, why are so many early deployments crashing and burning?

Emily That's the big.

Jeff Question, right? I mean, most failures aren't even due to bad tech, right? It's ignoring the demands of these high stakes regulated environments.

Emily Precisely. It's the friction between powerful tech and the real world. The source highlights three fatal traps.

Jeff Let's get into those.

Emily Trap one is giving the AI too much latitude.

Jeff Okay, what does that mean?

Emily Well, if you use an open ended prompt, like review this entire profile and see if it's bad, the AI will make plausible sounding but entirely baseless assumptions just to fill in the gaps.

Jeff Oh. Like hallucinations.

Emily Exactly. So teams end up spending more time validating the AI's hallucinated evidence than they would have doing the actual work.

Jeff That sounds exhausting. What's the second trap?

Emily Trap two is building for autonomy instead of trust. In other industries, minimizing human involvement is the goal, right?

Jeff Wait, actually, isn't the whole point of AI to be autonomous and save us time? Like, why wouldn't we want it to just run on autopilot.

Emily Because of the strict regulatory consequences in financial crime, a human must be accountable for every single suspicious activity report or SAR filed or every rule deployed. You can't just blame an algorithm if you unlawfully freeze someone's account. Ah, okay, that makes sense. The government doesn't care that your robot did it.

Jeff Exactly.

Emily Which leads to trap three explainability as an afterthought. You can't just slap a summary on a decision after the fact. That's not true. Explainability, right? Regulators need to see the exact mathematical signals used at every single step of the process.

Jeff So how do we actually design for trust then? Because skepticism from these fraud teams isn't just varied. It's well, it's required.

Emily It is absolutely required.

Jeff The real customers are affected by bad rules, and missing a signal creates massive regulatory exposure for the bank.

Emily Exactly. And in AML and fraud, explainability isn't just some neat feature. It is a baseline requirement for survival. Agents must show their work before deployment. They have to show the thresholds, the triggering criteria, everything.

Jeff It's exactly like high school math class. How so you know, it doesn't matter if you got the right answer. If you can't show the teacher exactly how you got there.

Emily That is a perfect analogy. Yes. And that's why automation without human oversight is so dangerous. Rules have to be drafted, then reviewed by a human, then deployed. SARS are drafted, then refined.

Jeff The human has to stay in the loop.

Emily Always. Plus auditability. Every action, every override, every single note must be logged automatically because regulators expect to see a completely pristine timeline.

Jeff Okay, so now that we know the pitfalls and you know the rules of the game here, let's look at the specific solution Data Visor's Conversational agent Vera.

Emily Yes. Vera.

Jeff And the big problem it solves in the current market is this idea of agent squads.

Emily Right. The fragmented approach.

Jeff Yeah. You have one bot for summaries. You buy another bot for building rules. And it creates this exhausting integration tax of constant handoffs between systems.

Emily It's terribly inefficient. But Vera represents unified execution a singular intelligence handling the whole lifecycle from start to finish.

Jeff Okay, let's unpack this because a big part of how it does that is through unsupervised machine learning or UML. Right.

Emily Yes, UML is critical here. It catches the unknown unknowns.

Jeff Wait. Meeting attacks we haven't seen before.

Emily Exactly. Coordinated attacks with no prior labels or reported losses. It detects them in real time just by seeing anomalous structural patterns forming.

Jeff That is wild. But if it's finding all these new threats, how do you keep the system from just getting clogged up with a million new rules?

Emily That's where technical hygiene comes in. Vera uses a search first protocol.

Jeff Oh, I read about this. So it searches the existing libraries before building anything new.

Emily Right. To prevent what we call signal sprawl in redundant rules. It keeps the whole system incredibly clean.

Jeff Which is so smart because signal sprawl just sounds frustrating. It's like, it's like having five different people buying the exact same groceries for one house. You just end up with way too much milk and no room in the fridge.

Emily Exactly. The search first protocol stops that technical debt before it even starts, and then you have strategy agility instead of waiting days or weeks for engineers to build a new feature, analysts just use plain language. Say they need to track ACH, sum accumulations. They just ask Vera and it builds the missing feature instantly.

Jeff Okay, so we've built the rules. We caught the bad guys. We have no redundant groceries right now. How does this unified system handle the messy aftermath? Like the optimization and all that regulatory paperwork?

Emily It's actually really elegant for optimization. Vera tests parameters against historical data. It provides F1 scores, precision and recall metrics before committing to any changes.

Jeff So it proves it works first.

Emily Yes, and it gives natural language explanations for its suggestions. It will say "If you adjust this threshold, precision drops by this much, but recall goes up."

Jeff Wow, that's incredibly helpful. And I saw that it also helps with investigation efficiency. Right.

Emily It does. It ingests manual review comments directly from the frontline investigators to constantly tune the rules. Oh, nice. It completely bridges that gap between the high level strategy and the actual frontline workers.

Jeff That makes total sense. But what about the compliance side? The SARS, because I know that is a massive headache for banks.

Emily Oh it's huge. Vera aggregates all the case intelligence across linked profiles and actually generates comprehensive SAR narratives automatically.

Jeff Okay. But if the AI drafts this huge regulatory narrative for you, isn't there a huge risk the human just gets lazy and blindly hits approve?

Emily You'd think so. But no. This is where the conversational aspect shines. The human remains firmly in the loop. They literally chat with the draft to refine it.

Jeff Also, they interrogate it.

Emily Exactly. They act more like a senior editor. They tell the agent to expand on a specific wire transfer or fix a timeline. So they are actively shaping it, ensuring the bank can confidently stand behind the filing.

Jeff That is brilliant. And the source mentioned that all these capabilities are out of the box, right? Like no extra technical integrations beyond the standard setup.

Emily Correct.

Emily It eliminates those multi-year implementation nightmares that banks usually face.

Jeff Incredible. So to briefly recap for you listening, True, Agentic AI isn't some hands off magic wand you just plug into the wall?

Emily No. Definitely not.

Jeff It's an interactive, unified collaborator. It turns days of manual fraud hunting into literally minutes of plain language execution, all while keeping a meticulous, regulator friendly paper trail.

Emily It fundamentally shifts the posture of a bank from reactive to proactive.

Jeff It does, which leaves us with a final, really provocative thought to ponder.

Emily Oh, I know where you're going with this.

Jeff Yeah. I mean, if systems like Vera are using unsupervised machine learning to instantly map out the unknown unknowns of criminal behavior, right. What happens in a few years when the criminals' AI and the banks' AI are locked in this invisible real time war? It's just constantly adapting to each other at computational speeds that no human can even track.

Emily Well, basically just be the governors of the machines at that point.

Jeff Yeah, it's wild to think about you definitely won't be bringing a manual spreadsheet to that fight.

Emily Definitely not.

Jeff Well, that's all the time we have for today. Keep diving deep, everyone.

Jeff You've been listening to "*What the F Happened? Fraud and Financial Crime Deconstructed,*" a DEFEND podcast by DataVisor. If you want to keep learning between episodes, check out DEFEND Training, a set of self-paced online courses for fraud and financial crime professionals, practical and built around real world scenarios.

Emily And you can earn CPE credits through the ACFE San Francisco Bay area chapter. You can find it at datavisor.com/defend- training. The link is in the description.

Jeff This episode's audio was generated using Google's Notebook LM based on expert analysis and trusted sources. Thanks for listening. We'll see you next time.