EPISODE 30 | Monthly Brief

AI Fraud War

August 3, 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 So in the I'd say roughly three hundred milliseconds. It just took you to blink. An AI agent you have literally never heard of, just decided whether to freeze your entire life savings or let a hacker completely empty your account. Yeah, just like that.

Emily It didn't ask a human manager for permission. It didn't, you know, send an email. It analyzed thousands of data points, cross-referenced your device ID with global networks, and made a legally binding financial decision before your eyelid even finished opening.

Jeff I mean, it is a profound shift. We are really moving away from an era where banking security was defined by a human analysts just looking at spreadsheets, right? And we're stepping into this totally new arena where machines are waging a silent, millisecond by millisecond war over our digital trust.

Emily And that is exactly our mission. Today, we are exploring how the entire financial system is rapidly and rapidly rewiring its operating model to fight back against a completely new era of AI driven financial crime.

Jeff Which is so necessary right now. Absolutely. And our insights today are pulled from the August Defend Newsletter by DataVisor, which is an award winning fraud and risk platform operating right on the front lines of this stuff.

Emily They really are.

Jeff And to set the stage for you, there is a quote in this newsletter from DataVisor CEO Patrick Harr. Okay, let's unpack this. He notes that because fraud is being fundamentally reshaped by AI, financial institutions need a system that responds with the exact same speed and intelligence as the threats they face.

Emily And that, quote, it really gets straight to the central friction in modern finance. I mean, if the attackers are using generative AI and automation to move at literal machine speed.

Jeff Right, which.

Emily They are exactly defending your institution with human speed is just, well, it's a guaranteed way to lose. It's a mathematical certainty. You just cannot outthink an algorithm using manual processes.

Jeff Which means the system has to adapt and fast. So today we're going to break down the changing laws that are finally trying to keep up the what they call shape shifting nature of modern fraud, and ultimately, how banks are building literal fleets of AI agents to protect your money.

Emily It's fascinating. But, you know, before we even touch the high tech AI orchestration, we actually have to start with the rule book, right?

Jeff The legal side.

Emily Yeah. Because you cannot fight fast moving AI if banking regulations legally force you to move slowly or, you know, actively penalize you for taking preventative action, right?

Jeff I mean, for decades, the law has essentially punished banks for being overly cautious. Think about it from the consumer side. Like when you transfer money, you want it instantly.

Emily Oh for sure.

Jeff If a bank holds your money for three days to investigate, you get furious. So we had these regulations that mandated extreme liquidity in speed.

Emily Exactly. The historical burden has always been on making funds available immediately. But if we connect this to the bigger picture, regulators are finally realizing that forcing instant liquidity in an era of instant AI driven fraud is basically handing the criminals a getaway car with the keys already in the ignition.

Jeff Wow.

Emily Yeah. Which is why we are seeing two massive policy shifts that are completely redrawing the legal battle lines right now.

Jeff So let's start with the first one, which is a major clarification from the OCC and the FDIC regarding the three fourteen B safe harbor. And I know a lot of risk professionals have been waiting for this exact clarification for years.

Emily Oh, they absolutely have. So traditionally, section three, fourteen B allowed financial institutions to share information with each other, but it was incredibly restrictive.

Jeff Like they couldn't just chat about anything, right?

Emily They generally had to establish that they were specifically combating money laundering or terrorist financing. The threshold to share data was incredibly high.

Jeff So if a bank saw a transaction that looked like, you know, an account takeover or a scam, but they couldn't definitively prove the money was being laundered, their hands were just tied.

Emily Exactly tied. Privacy laws and the fear of litigation kept them totally siloed.

Jeff But this new OCC and FDIC clarification changes that threshold for matically.

Emily They just confirm that these safe harbor protections extend to suspected fraud as well. A reasonable suspicion of fraud is now enough to trigger the safe harbor.

Jeff That's a huge pivot.

Emily It is. And the technical scope of what they can share has just exploded. They can share cyber, physical and behavioral signals in real time, pre-settlement and even without a prior relationship.

Jeff Wait. Explain the mechanics of that pre-settlement sharing because that seems like the real game changer here.

Emily Okay, let's say a bad actor initiates a wire transfer from Bank A to Bank B. Historically, Bank A might not realize it's fraud until the money has already settled in Bank B, right?

Jeff At which point the fraudsters already transferred it to a crypto exchange or something.

Emily Exactly. Now, Pre-settlement Bank A's system can instantly ping Bank B through an API and say, hey, the IP address initiating this transfer matches a known bad IP. The geolocation is suddenly three thousand miles away from the customer's home, and the device ID is brand new.

Jeff Wow. So they are passing telemetry data instantly. It's almost like it's like giving neighborhood watch groups, walkie talkies, and the legal right to lock the neighborhood gates before the getaway car even leaves

the subdivision, rather than just forcing them to quietly take notes for the police after the house has already been robbed.

Emily That is exactly the shift in posture, moving from reactive forensics to proactive preventative action.

Jeff Wait, I have to push back here.

Emily Sure.

Jeff If banks are now legally encouraged to share my geolocation, my specific device ID, and hold my funds instantly based on just suspicion. Doesn't this mean my perfectly legitimate transfer is constantly at risk of being frozen?

Emily That's a fair point.

Jeff Like if I log in from a new phone on a beach in Mexico and try to pay my mortgage, am I just going to get locked out of my own life because some algorithm got suspicious?

Emily And that is the exact tension regulators are trying to navigate. And it leads us directly to the second major policy shift, which is the Stop Payments Fraud Act, or H.R. nine hundred thirty one. This amends regulation CC, also known as the Expedited Funds Availability Act.

Jeff Write, read. CC is the law that historically mandated that instant liquidity we just talked about.

Emily Yes, exactly. The Stop act amends it to give financial institutions explicit legal cover to place temporary risk based holds on suspicious wires, checks and instant transfers without facing massive liability for delaying a payment.

Jeff Okay, but what about the beach in Mexico?

Emily Right. So to address your concern, the Stop act requires these holds to be incredibly objective and rigorously documented.

Jeff So risk teams can't just apply broad, indiscriminate holds. They can't just flip a switch and say, all international transfers get paused today.

Emily No, not at all. They have to point to specific documented criteria. Like what? Like a distinct mismatch between your biometric behavior and your historical login patterns, or a known scam signature, or maybe a sudden unexplained routing to a high risk jurisdiction. Got it. This ensures banks stay compliant with fairness and anti-discrimination rules while still getting the operational breathing room. They need to actually coordinate with law enforcement and their three hundred and fourteen B partners before the money becomes irreversible.

Jeff Okay, so the regulatory landscape has adapted. The law is finally letting banks talk to each other and pause suspicious transactions without getting penalized. But here is the catch. Even with this new explicit legal cover, banks can't actually utilize it if they are relying on human analysts.

Emily No, they really can't.

Jeff The legal window is open, but the speed of the threat has completely outpaced human cognition completely.

Emily And that is why the operating model is changing. The threats have fundamentally evolved, and the window to stop them has literally shrunk to absolute milliseconds.

Jeff And a lot of this is driven by what the DataVisor newsletter calls polymorphic threats. Let's dig into the mechanics of this, because I think people hear AI fraud and just think of, you know, deep fake voices or something. Polymorphic fraud is something entirely different.

Emily It is. So in the past, fraud detection relied on a rules based system. A fraudster uses a specific IP address or a specific sequencing of API calls, and the bank identifies the signature of that attack. A human engineer then writes a rule into the system that basically says, if you see this signature, block it.

Jeff But with polymorphic AI fraud, there is no static signature. It mutates.

Emily Exactly. AI generated fraud never attacks the exact same way twice. It might use synthetic identities where the parameters, the timing of the keystrokes, the routing paths, the specific device telemetry, all of that is slightly randomized every single time it attempts a login or a transaction.

Jeff That sounds impossible to track.

Emily It is if you're using old methods, there is a structural ceiling to relying on recognizing known attacks. You cannot build a defense based on looking in the rear view mirror if the threat is constantly generating a new disguise.

Jeff Which means financial institutions need a framework for detecting mathematical anomalies that have absolutely no prior label. and they have to do it with the speed that is almost impossible to comprehend. And here's where it gets really interesting. The newsletter details this one hundred millisecond budget.

Emily Yes.

Jeff Accurate detection of a polymorphic threat is entirely useless if the decision to block the transfer arrives after the money has already cleared. So real time decisioning has to happen inside a one hundred millisecond latency budget.

Emily It is a brutally unforgiving time frame, right?

Jeff I mean, to put that in perspective for you, listening human reaction time to visual stimuli is roughly two hundred and fifty milliseconds. So by the time your brain even registers that you are looking at a transaction on a screen, the millisecond budget has already expired and the money is gone and the fraudster has vanished.

Emily What's fascinating here is how DataVisor approaches this to pull off a defense. Within that budget, fraud teams have to meticulously measure three distinct simultaneous clocks the decision clock, the adaptation clock, and the response clock.

Jeff Let's walk through those mechanisms. So the decision clock is pretty straightforward. That's the initial one hundred millisecond window to analyze the inbound transaction and decide if it's fraudulent.

Emily Yes, but the adaptation clock is where the real AI battle happens. This is how quickly the system can learn from new unlabeled data and update its own internal models.

Jeff Okay, so if a polymorphic attack starts probing a network with the brand new mutated strategy.

Emily The adaptation clock measures how fast the unsupervised machine learning models can identify the mathematical variants of that new attack, and autonomously adjust its detection parameters without a human engineer ever writing a new rule.

Jeff Because if a human has to identify the mutation, write a rule, test the rule, and deploy it. You're talking about weeks of adaptation time.

Emily Exactly. The AI has to rewrite its own defense in real time.

Jeff Wow. And then finally, you have the response clock, right?

Emily Once the AI decides something is anomalous, how fast can the system execute the complex response? Freezing the funds, leveraging the three fourteen B safe harbor to instantly alert the pure bank's API and locking down the user's account.

Jeff And this latency budget applies to the entire compliance stack, not just fraud. Right. Because the newsletter explicitly points out the impact on sanctions screening.

Emily It does. And this is a huge operational shift. Historically, checking customers against P politically exposed persons and adverse media lists was often done as an overnight batch job.

Jeff Like they'd just wait until the end of the day.

Emily Yeah, the bank would gather all the transactions, send them to a provider like LexisNexis, and review the flags the next morning.

Jeff But sanctions risk doesn't wait for batch processing anymore. If instant payment rails are moving money to foreign jurisdictions in seconds, you can't afford to find out tomorrow that you just facilitated a transfer for sanctioned entity, right?

Emily Which means keeping those three clocks under one hundred milliseconds requires unifying fraud signals and AML anti-money laundering into a single workflow. It's called the frame a life cycle. The AML. Yes, you can no longer afford the technical latency of bouncing data back and forth between siloed fraud systems and siloed AML systems.

Jeff Okay, but this exposes a massive logical paradox in the industry right now.

Emily What do you mean?

Jeff Well, we have these powerful new regulatory tools like the Stop act, but they require objective, documented criteria to use. We have polymorphic threats mutating instantly, requiring sub one hundred millisecond decisions. Human teams physically cannot spot the threat. They cannot manually coordinate with other banks. And they certainly can't compile complex legal documentation in the tenth of a second.

Emily No they can't.

Jeff So how do banks actually execute on this?

Emily They execute by fundamentally changing what a bank is. They are moving away from being organizations of humans, assisted by software and becoming organizations of agentic AI overseen by humans Agentic AI.

Jeff Let's break down how this actually works across that framo lifecycle, because we are just talking about a simple chatbot analyzing a spreadsheet here.

Emily Far from it. An agent in this architectural context is an autonomous AI system that has agency. It doesn't just analyze data. It is authorized to take intelligent action based on that data within its specialized domain. For example, DataVisor highlights their AI reporting agent.

Jeff Right, and the newsletter notes this agent automates the generation of SARS and Ctrs, making compliance ready, reporting up to ninety percent faster, which.

Emily Is incredible.

Jeff And the mechanism here is crucial. The AI isn't just filling out a form, it is pulling the forensic data from the detection agent, mapping out the behavioral anomaly, and structuring the narrative for the legal filing with a complete audit trail that actually satisfies regulators.

Emily It is turning raw, high speed data into a documented, compliance ready narrative.

Jeff Which is how you satisfy the stop requirement for documented objective criteria without slowing down the defense.

Emily Exactly. The AI does the heavy lifting of documentation instantly, and this multi-agent approach is gaining massive traction. Datavisor recently earned three major industry awards the AI Breakthrough Award, Paytech Awards, and recognition from silent.

Jeff Oh, wow.

Emily Yeah, it really validates that this AI native unified platform is the necessary path forward.

Jeff And they're actually doing a deep dive into the mechanics of this on a webinar coming up on Wednesday, August twenty sixth. It's titled Trust at Scale Multi-Agent Orchestration featuring pix Wan Wang and Alicia

Govan. That's right. They're going to get into exactly what happens when detection agents, investigation agents and reporting agents hand off work to one another.

Emily It is going to be a critical conversation because that handoff process is where the true operational risk lies today.

Jeff Wait, I have to stop you there. Let's confront the paradox I mentioned earlier. So what does this all mean? If we have AI agents detecting the crime, AI agents investigating the crime, and AI agents writing the legal reports, how is a human actually involved?

Emily That's a great question.

Jeff The newsletter talks about human in the loop approvals, but we just established the human reaction time is way too slow for one hundred millisecond decision. How can a human be in the loop if the loop closes before they can even blink. That sounds physically impossible.

Emily This raises an important question, and it requires clarifying what in the loop actually means in a multi-agent architecture. The human isn't sitting there clicking approve on every single microtransaction that would defeat the entire purpose.

Jeff So what are they doing? Are they just drinking coffee?

Emily No. They are governing the risk thresholds and auditing the complex edge cases. Think of it this way. The AI handles the one hundred millisecond autonomous actions, freezing the funds, blocking the IP, compiling the initial report based on strict parameters set by the human team. Okay, the human is in the loop post action for the review process or they are engaged. When an investigation agent encounters a highly nuanced scenario that falls outside its confidence threshold. The human manages the strategy. The agent executes the tactics.

Jeff Okay, that makes more sense. The human is designing the rules of engagement, and the AI is fighting the battle at machine speed.

Emily Exactly.

Jeff But that brings us to the real danger discussed in the upcoming webinar, which isn't the AI like turning evil and going rogue. The real threat to institutions right now is something called orchestration, sprawl, orchestration.

Emily Sprawl is the silent killer of AI adoption in finance.

Jeff How so?

Emily Well, when a bank scales from one successful pilot agent to a massive fleet of specialized agents running in production, the complexity multiplies exponentially.

Jeff It's like an operating room where you have ten of the world's most brilliant robotic surgeons, but no head doctor calling the shots.

Emily That's a perfect analogy.

Jeff Like the detection agent is cutting, the reporting agent is suturing, the sanctions agent is monitoring vitals. But because they aren't communicating with perfect synchronization, they start interfering with each other. The patient bleeds out on the table, not because the robots lacked skill, but because they lacked orchestration.

Emily Spot on. If you have a dozen disconnected AI agents running across the life cycle without proper role based controls and cross agent audit trails. You create dangerous new blind spots.

Jeff So instead of giving you leverage against polymorphic threats, a poorly orchestrated fleet just creates a chaotic internal environment, right?

Emily Threats slip through the cracks of the system's own complexity.

Jeff Which defeats the purpose.

Emily Exactly. Proper multi-agent orchestration means meticulously managing the life cycle of these agents handling rollbacks. If an adaptation clock updates a model incorrectly, and ensuring that the data handoffs between the detection agent and the reporting agent are seamless and mathematically auditable.

Jeff Wow. We have covered a tremendous amount of ground today, and it is just fascinating how all of these pieces interlock. Let's just briefly recap this journey for you, because it really is a complete rewiring of the financial system.

Emily There really is.

Jeff It starts with regulators finally updating the rule book through the three fourteen B clarification and the Stop act. Banks finally have the explicit legal cover to hold suspicious funds and share telemetry data in real time without being penalized for moving slowly.

Emily But those legal powers are completely useless if the bank relies on human speed because they are facing polymorphic fraud. Those AI generated attacks that mutate their parameters to avoid static rules.

Jeff Which means banks have to operate inside a hundred millisecond latency budget, meticulously measuring their decision adaptation and response clocks. And since humans literally cannot perceive data that fast, the only viable solution is deploying multi-agent AI fleets across the entire fraud and AML lifecycle. AI agents detecting, investigating, and generating compliance reports instantly while humans manage the orchestration to prevent that deadly sprawl.

Emily It is a complete paradigm shift in the mechanics of global finance.

Jeff Absolutely. And as we wrap up this deep dive, I want to hand it over to you for a final thought. We've talked a lot today about these fleets of agents handing off tasks, communicating via API, executing decisions in milliseconds. If you project this out, where does this leave us?

Emily It leaves us with a really profound structural oppression about the very nature of a bank.

Jeff I'm listening.

Emily Well, if we are rapidly moving toward a reality where your bank's internal AI agents are constantly handing off work to other AI agents across the entire life cycle of a transaction, detecting anomalies, investigating histories, instantly communicating with pure banks through three hundred and fourteen B safe harbors and writing legal reports, all in a fraction of a second. At what point does the concept of a human led bank cease to exist? We have to ask ourselves, are are we moving toward a future where a bank is simply a massive server farm of specialized AI agents, legally negotiating trust and risk with other banks, AI agents at the speed of light?

Jeff Why are we banking with institutions? Or are we just plugging our money into a self-governing, global AI network? That is definitely something to think about. The next time you tap your phone to pay for coffee and wonder what exactly happened in those three hundred milliseconds. Thank you so much for joining us on this deep dive. Keep questioning the systems operating beneath your everyday digital life and we will catch you next time.

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.