EPISODE 34 | Analysis

The Race to Outsmart Financial Crime

AUGUST 31, 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 Okay, so let's unpack this. All right. I want you to imagine a massive logistics network, uh, a delivery company that moves roughly six billion packages every single year.

Emily Which is a staggering volume.

Jeff Right? It's massive. Now, imagine this kind of invisible army of automated bots just quietly infiltrating that network. And they aren't hijacking trucks or anything like that. They're actually creating thousands upon thousands of fake accounts in literally a matter of seconds.

Emily Yeah. Using synthetic identities.

Jeff Exactly. And they use those identities to seamlessly reroute flat screen TVs, uh, high end electronics while they're already in transit.

Emily Wow.

Jeff So by the time the legitimate customer is, you know, wondering where their package is, the thief has just completely vanished.

Emily Yeah. Long gone.

Jeff So the question is, how do you stop a thief who can't even see operating at a speed that human beings basically can't even process?

Emily Well, it is. It's really the defining security challenge of the modern financial ecosystem right now, because we are dealing with these highly coordinated syndicates and they operate like, well, like well-funded tech startups. Really?

Jeff Oh, absolutely.

Emily They're leveraging all this automation just to overwhelm our legacy defense systems.

Jeff And that is exactly our mission for this deep dive. Today, we're exploring how artificial intelligence is. Well, it's finally moving past all the flashy marketing hype, right?

Emily Finally.

Jeff And it's fundamentally rewiring how institutions fight financial crime at a massive scale. Yeah. So we are pulling our insights today from a brilliant stack of research. Yeah. Uh, we've got the Q3 2026 Forrester Wave report on financial crime management solutions.

Emily A very thorough report.

Jeff It really is. And we're pairing that alongside a detailed companion perspective piece by DataVisor, which analyzes those exact findings.

Emily Yeah, the DataVisor piece provides some great context.

Jeff It does. So we're going to look at why the traditional, you know, these heavily guarded walls between fraud prevention and anti-money laundering are finally collapsing.

Emily Which is huge.

Jeff Yeah, and what it actually looks like when true AI is put to work against billions of real world transactions.

Emily You know, what's fascinating here is the distinction the Forrester Wave report makes regarding the the weapons we actually use to fight this battle. Okay. Specifically the evolution of the AI itself, because a lot of what currently exists in the industry right now, uh, even when vendors brand it as an AI agent, right?

Jeff The buzzwords.

Emily Exactly. It's essentially just a sophisticated batch job.

Jeff Wait a batch job. When I hear the term batch job, my mind immediately goes back to like, legacy computing.

Emily Yeah, yeah.

Jeff Like feeding a system, a massive pile of data, letting it churn overnight and just, you know, hoping for a useful printout in the morning.

Emily Well, sadly, that is functionally how a lot of these early AI tools operate within financial crime management today. Really? Yeah. They run very passively behind the scenes. So the system might collect a risk score, maybe scrape together some basic evidence and reason codes. Okay. And then it uses basic natural language processing to write up a static narrative that summarizes an alert.

Jeff Just a write up.

Emily Yeah. And then it hands that digital paperwork over to a human investigator. I mean, sure, it removes a bit of administrative friction, but it is entirely reactive. The AI isn't really investigating.

Jeff It's just taking notes.

Emily Exactly. It's just taking notes.

Jeff So if we use like a library analogy here, a basic AI chat bot is maybe like a librarian pointing you to a specific aisle. Like here's the data. Go read it. Right. And the batch job AI you're describing is maybe one step up. It's like a research assistant who reads the book, highlights a few passages, hands you a summary, and then just clocks out for the day.

Emily That is a perfect way to put it.

Jeff But that's not enough to stop the invisible bot army. We talked about at the top of the show. I mean, it can't be.

Emily Not even close. What the Forrester Report really highlights is the critical need for interactive, conversational co-pilots. Okay, true AI agents that dynamically participate in the investigation as it's happening. Yeah. And the report notes that these are actually still quite rare in the market. Most vendors just haven't perfected that capability yet.

Jeff Right. It's harder than it sounds on a marketing brochure.

Emily Much harder. However, Forrester specifically points to DataVisor and their agent, which is named Vera, as the prime example of an AI that helps human teams translate abstract ideas into immediate platform execution.

Jeff Oh, wow.

Emily Yeah. In fact, on the specific criterion of AI agent use in the solution, DataVisor scored a 5.0.

Jeff Which is the highest score, right?

Emily The highest possible score you can get. Yes.

Jeff So to elevate that earlier library analogy, a true AI agent isn't just a research assistant handing you a summary. It functions a lot more like a central nervous system for the whole platform.

Emily I like that.

Jeff Yeah. Like it doesn't just log the anomaly. It actively synthesizes the evidence across the entire network, and then it guides the human investigator on exactly where to look next.

Emily Yes. And it goes further. It actually helps rewrite the detection logic to permanently immunize the system against that specific type of attack.

Jeff Wow.

Emily Okay. That underlying mechanism, you know, the ability to act and adapt on the fly, that's what separates a true agent from just a chat bot, right?

Jeff Because financial crime analysts aren't just working from a single simple prompt.

Emily No, not at all. They have to evaluate these incredibly complex, obscured relationships. We're talking across multiple accounts, varying IP addresses.

Jeff Shell companies hit the web.

Emily It is. And the AI agent has to navigate that whole workflow right alongside them. And crucially, it must do this while providing absolute defensibility defensibility.

Jeff Yes, the source is hammered on that concept quite a bit and it makes perfect sense to me.

Emily Oh yeah.

Jeff Because when a bank freezes someone's account or blocks a massive corporate wire transfer, The compliance officer can't just tell a federal regulator. Well, the black box algorithm said it was a bad transaction.

Emily No, the regulatory landscape simply does not allow for black box decision making. You'd be fined into oblivion, right? An AI agent must provide concrete reason codes for all of its automated decisions. It has.

Jeff To show.

Emily Its work. Exactly. It has to show its mathematical homework, basically explaining the precise logic it used to flag a transaction so a human analyst can actually audit and verify that logic.

Jeff Which brings us to the ultimate proving ground, really, because having a brilliant, auditable AI nervous system, that's a beautiful concept in a sterile testing environment. Sure. But the true test is production out in the wild.

Emily Always. If we connect this to the bigger picture, the performance bar just fundamentally changes when an AI agent is handling live data streams. Right? Which brings us back to that delivery services company. You mentioned earlier.

Jeff The six billion packages.

Emily six billion packages annually. That level of scale requires processing millions of data points every single second.

Jeff That's hard to even wrap your head around.

Emily It is. And the mass registration attacks they were facing are notoriously difficult to stop.

Jeff So here's where it gets really interesting to me. The strategic goal for them wasn't just catching the fraud. It is entirely about the timing of the catch.

Emily Yes, the timing is everything.

Jeff Because mathematically, catching a fake account after it has already rerouted a package, that's just a permanent financial loss. You're already out the money.

Emily Yeah, the damage is done.

Jeff So the mandate is to catch the fraudulent account at the exact millisecond of signup. But mechanically speaking, why is point of signup fraud so exponentially harder to detect than catching a criminal mid theft?

Emily It really just comes down to a severe information deficit. Okay, think about it. When a brand new account is created, you have zero behavioral history to analyze.

Jeff Oh, right. Because they just got there.

Emily Exactly. Yeah. You don't know their normal purchasing habits. You don't know their average transaction size or even what time of day they usually log in. Yeah, the only data available is what they present to you in that exact moment. So an IP address, a device ID, an email, maybe a phone number.

Jeff And these sophisticated syndicates, they aren't logging in from known malicious servers that are easy to flag.

Emily No, they use very clean, localized IP addresses, pristine newly generated emails that haven't been blacklisted yet.

Jeff Right. So they just blend right into the noise of normal, legitimate customer traffic.

Emily Exactly. So to bridge that information deficit at the exact millisecond of sign up, the AI system cannot rely on historical profiling.

Jeff Because there is no history.

Emily Right? It must instantly normalize and analyze massive high velocity data sets across the entire global network, normalizing.

Jeff So like standardizing the data.

Emily Yeah. By normalizing, we mean the AI takes a device ping formatted in one coding language, an IP address in another format, and it standardizes them instantly so it can run them against a unified risk model. Wow. The AI has to recognize that this specific device ID, even though it looks perfectly clean in isolation, it actually exhibits subtle micro behavioral connections to a broader cluster of new accounts that are acting collectively suspicious.

Jeff Okay, so it's looking at the velocity of the data, maybe the typing cadence, these really subtle interconnections that a human couldn't possibly map out in real time.

Emily No human could do it. That rapid integration of high velocity data is the absolute key. And by deploying a modern fraud platform that is capable of handling that instantaneous load, this specific delivery company achieved a sixty percent detection uplift sixty percent.

Jeff That is a staggering jump in threat neutralization.

Emily It's massive. It fundamentally altered their overall risk profile. And because the system analyzed the data holistically, rather than using blunt, rigid rules, they also saw a forty percent improvement in customer reviews.

Jeff Oh wow.

Emily Yeah, because they removed all the friction for legitimate users, and ultimately that saved them four million dollars in direct fraud losses.

Jeff That is incredible. Okay, so we have seen how AI neutralizes external threats, right? The bot armies, the mass registration, right. But to truly understand the current state of financial crime, we really have to look inside the house because historically, one of the biggest vulnerabilities for financial institutions hasn't been external hackers outsmarting the system. No, it has been a legacy structural flaw in how banks organize their own internal defenses.

Emily The internal silos.

Jeff Yes. The silos. I was thinking about how to visualize this. It is like a bank hiring two highly elite security guards to protect the exact same building, but they strictly forbid them from using their walkie talkies to communicate.

Emily That is a great way to look at it.

Jeff Right? So one guard stands at the front doors just scanning for pickpockets and fast moving shoplifters. That's your fraud prevention team operating in real time. But the other guard is stationed way down in the basement vault, meticulously looking for long term inside jobs and organized syndicates, slowly stashing money that is your anti-money laundering or AML team. Yeah, they are defending the exact same ecosystem, but their data never crosses paths.

Emily Never. And that fragmented reality has plagued compliance teams for decades now. I could imagine they run these completely parallel software systems that are fundamentally incapable of speaking to one another. But the Forrester report notes a massive structural shift happening in the industry right now. Okay, vendors are finally breaking down these walls by sharing machine learning models and risk scoring telemetry across both disciplines. The industry term for this convergence is fram f r a ML fraud and anti-money laundering framework.

Jeff Okay, I get the theory behind combining fraud and AML into a unified framework platform. I mean, criminals commit fraud to get the money and then they have to launder it to actually use it.

Emily Exactly. It's one continuous criminal life cycle.

Jeff Right. But aren't these departments heavily siloed for a valid practical reason? I mean, fraud relies on split second real time blocking, whereas AML is typically this slow, methodical investigation that can take months.

Emily Yeah.

Jeff Does it? Merging those two wildly different operational tempos just create a massive technical and regulatory headache for a bank.

Emily Well, you've hit on the exact friction point that prevented this integration for so long. The operational tempos are very different, but the underlying data signals required to catch both crimes are practically identical.

Jeff Oh, interesting.

Emily Yeah. The breakthrough isn't about forcing AML investigators to suddenly work at millisecond speeds. It is about allowing the AML models to seamlessly ingest the real time behavioral anomalies that the fraud side is already catching, and the software tooling is finally mature enough to support both of those workflows simultaneously. The DataVisor Perspective Brief highlights a phenomenal case study to illustrate this exact thing.

Jeff This is the NASA Federal Credit Union case.

Emily Yes. NASA Federal Credit Union, which is an institution managing over three billion dollars in assets.

Jeff Okay. So massive scale again. And prior to their upgrade, they were suffering from the exact siloed infrastructure we just described separate systems, duplicated workflows, separate data pools, completely disconnected. So if an investigator was looking into a suspicious wire transfer from a fraud perspective, they

had absolutely no visibility into whether a compliance officer literally down the hall was investigating that same customer for money laundering violations.

Emily Which creates a incredibly dangerous blind spot. Not to mention an extraordinary waste of human capital.

Jeff Oh for.

Emily Sure. So to solve this, NASA FCU migrated to a single unified framework platform. And the absolute linchpin of this whole transformation was their utilization of a knowledge graph tool.

Jeff A knowledge graph. Let's break down the mechanics of that because it sounds like abstract tech jargon, but from what I read, it's actually a highly visual, incredibly powerful AI tool.

Emily Oh it is, it really changes the game. Instead of a human investigator staring at, you know, endless static rows on a spreadsheet trying to find correlations in their head.

Jeff Which sounds awful.

Emily It is awful. The Knowledge Graph dynamically maps out this multi-dimensional web of relationships visually on the screen. It connects accounts, device IDs, physical addresses, transaction histories all in real time.

Jeff Kind of like the detective corkboard with the red string connecting all the pushpins.

Emily Yes, exactly like that. But digitized and instant.

Jeff Okay, so if a bad actor manages to spin up a synthetic identity, but they accidentally use the same physical mobile device that was previously associated with a known fraudulent account, like three years ago. Mhm. The knowledge graph just draws a bright red line connecting those two nodes instantly.

Emily It makes the invisible syndicates completely visible. And by surfacing those hidden linkages automatically, NASA FCU achieved remarkable operational metrics within just nine months of deploying the platform.

Jeff Nine months is fast.

Emily Very fast. In the enterprise world. They saved over three hundred manual hours of investigation time, and the system immediately flagged thirty six percent of all risky transactions without requiring any human intervention at all. Wow. Which directly led to a forty one percent drop in manual review time across the board.

Jeff And the administrative impact was even larger than that. They recorded a 90% drop in SAR filing time.

Emily Huge.

Jeff Just for those unfamiliar, SAR stands for Suspicious Activity Report. Banks are mandated by federal law to file these really comprehensive reports with the government whenever they spot potential money laundering or severe fraud.

Emily And the paperwork burden for that is immense, right?

Jeff So slashing that administrative time by ninety percent is a monumental victory for any compliance department. They calculated about two hundred thousand dollars in direct annual savings purely from operational efficiency.

Emily Now, those efficiency metrics are fantastic, but they introduce a concept that is deeply painful for anyone who actually works in risk management, right? And that is the hidden tax of the industry.

Jeff False positives.

Emily Yes.

Jeff Because catching criminals is obviously the objective. But if your detection system is so hypersensitive that the alarm goes off every time a stray cat walks past the building, your human security guards are going to just be completely overwhelmed by the noise.

Emily Exactly.

Jeff And the sources describe false positives as a literal tax on institutional resources. Highly paid analysts spend their days chasing fake alarms at the direct expense of investigating sophisticated real world threats.

Emily This raises such an essential point about the mechanics of system tuning, because, you know, it is remarkably easy to build a compliance system that stops one hundred percent of fraud. Really, you simply configure the software to decline every single transaction that attempts to process.

Jeff A perfectly secure vault that literally no one can ever open.

Emily Exactly. The true engineering challenge is precision. How do you cut the noise without missing the actual signal? With their Unified framework architecture, NASA FCU reduced their AML false positives by a staggering forty two percent.

Jeff That's massive.

Emily And they did that while simultaneously retaining one hundred percent of their true positive detections.

Jeff So to understand the how behind that, it really comes down to contextual awareness versus rigid rules, right? Yes. A legacy system basically operates on rigid thresholds. So for example, a rule might dictate that any study wire transfer over five thousand dollars immediately triggers an alarm.

Emily Which happens constantly.

Jeff Right. Which creates a massive false positive rate for perfectly normal business activity. But the AI agent evaluates the historical context of a user's behavior. So it recognizes that while a five thousand dollars transfer is highly anomalous for user A, user B routinely purchases expensive art on Friday afternoons, so the AI just approves it without even waking up a human investigator.

Emily That contextual intelligence is exactly why a forty two percent reduction is achievable, and I think it transitions perfectly into the specific warnings detailed in the Forrester Report.

Jeff Yes, the warnings.

Emily Because every vendor on the market promises that level of magic precision, but giving the right alert to the right person doesn't actually matter if the underlying system architecture is fundamentally flawed.

Jeff Okay, let's strip away the consultant speak here.

Emily Let's do.

Jeff It. Forrester explicitly warns enterprise buyers about Q and case management imbalances. What does that failure actually look like on the floor for an analyst dealing with a massive backlog of alerts?

Emily Well, it means that the glossy marketing brochure is often masking severe operational bottlenecks.

Jeff Okay.

Emily Forrester lays out a few specific technical warnings for buyers evaluating these financial crime management solutions. First, they emphasize that Q routing schemes.

Jeff That's the automated systems that distribute the alerts to the human workers.

Emily Yes, they warn that these schemes are wildly inconsistent across vendors. Some legacy systems use a very basic round robin distribution.

Jeff So just handing out complex money laundering alerts like a dealer passing out a deck of cards?

Emily Exactly. Completely ignoring the experience level of the investigator receiving the alert.

Jeff Oh it's terrible.

Emily It is a recipe for operational failure. If a junior analyst is randomly assigned a deeply complex multinational layering scheme simply because it was their turn in the queue, the investigation will stall out immediately.

Jeff Yeah, they wouldn't know what to do with it.

Emily So Forrester advises buyers to ensure the platform actually supports intelligent time based escalation or forced assignment routing. That's based on the specific risk parameters of the case and the historical performance of the analyst.

Jeff That makes total sense. Yeah. And the second warning focuses on the case management interface itself. Because if you purchase a platform that only offers a basic triage screen.

Emily Just an endless list of transactions.

Jeff Right, with a simple approve or decline button. Yeah, your human analysts are going to burn out instantly. The advanced vendors, they provide highly customizable workspaces where the AI agent actively pulls third party data directly onto the screen, cross-references public records, and automatically compiles the evidence package before the analyst even clicks into the alert.

Emily And that ties into the third warning Forrester issues, which is perhaps the most critical one regarding how these systems fundamentally understand the movement of capital. They warn that built in detection models must feature regulator driven, payment type specific support.

Jeff Meaning the system has to inherently recognize that not all money moves the same way.

Emily Precisely. If a vendor treats all financial movement uniformly, like applying the exact same behavioral logic to an international wire transfer as they do to an ACH transfer, the solution will fail.

Jeff An ACH transfers. For context, that's the Automated Clearinghouse protocol banks use to move money directly between domestic accounts. Right. Like the direct deposit or online bill pay?

Emily Yes, exactly. And ACH transfers have entirely different velocity and risk profiles than an instant, irrevocable wire transfer. The behavioral patterns of, say, a first party credit card scam do not look anything like the patterns of a highly coordinated corporate wire fraud attack. Forster is so adamant about this distinction that they explicitly recommend buyers obtain written confirmation of this payment type support before ever signing a contract.

Jeff Wow. When an industry report advises you to get a technical capability in writing, it is a glaring signal that the market is full of vendors who are just overpromising and under-delivering on their core architecture.

Emily It's the ultimate sign of a maturing market. It's forcing buyers to really look deeply under the hood.

Jeff Well, we have covered incredible ground today.

Emily We really have.

Jeff We analyzed the shift from passive batch jobs to proactive AI copilots. We explored how high velocity data normalization neutralizes mass bot attacks at the point of signup. We broke down the silos to understand Friemel, we visualize the knowledge graph and dissected Forrester's technical warnings to enterprise buyers. So stepping back, what does this all mean? Why should the average listener who doesn't run a corporate compliance department care about the underlying architecture of financial crime software?

Emily Because this invisible layer of technological infrastructure dictates the safety and the fluidity of your daily financial life.

Jeff Yeah, it is the structural difference between your entire life savings account being protected from a highly coordinated synthetic identity takeover. And on the flip side, it is the reason you're completely legitimate attempt to wire a down payment for a house gets frustratingly frozen by a rigid, antiquated algorithm on a Friday afternoon.

Emily Exactly. When global financial systems rely on outdated, siloed rules, the criminals exploit the blind spots, and legitimate consumers suffer the friction, right? But when institutions deploy unified framework models powered by dynamic AI agents, the defensive friction disappears for the consumer and the security net tightly closes around the syndicates.

Jeff So as we look forward, how does the Forrester Report see this AI integration evolving? We discussed AI moving deeper into the operational workflow, but what is the actual tangible roadmap here?

Emily The industry roadmap for these true AI agents is structured around five core pillars of capability.

Jeff Okay. Five pillars.

Emily First is assist. The AI actively explains anomalous activity. It summarizes complex historical data and surfaces the relevant evidence for the human.

Jeff Okay. That's the baseline central nervous system we discussed earlier.

Emily Exactly. Second is investigate. Here the agent proactively follows behavioral leads across multiple devices and external entities. It leverages the knowledge graph we talked about to map out relationships, and it recommends which subtle connections deserve immediate human attention. Very cool. Then third is build.

Jeff Wait. Build. Yeah. As in the AI itself is actually writing the detection code. Yes. That is a massive paradigm shift from just analyzing data.

Emily It is the most critical leap forward in the space. The agent acts as a direct translation layer. So it takes an analyst strategic intent. Let's say the analyst spots a brand new highly specific routing scam. The AI automatically writes the complex software rules the features and the detection logic necessary to catch that exact scam across the entire global network moving forward.

Jeff So it effectively removes the massive bottleneck of waiting weeks or months for human software engineers to update the legacy code.

Emily Exactly. You bypass that bottleneck entirely, which leads directly into the fourth Pillar Act. Act. The AI transitions from making passive recommendations to initiating orchestrated actions within the bank's governed workflows. It autonomously freezes the compromised account, suspends the suspicious device ID and halts the wire transfer in real time.

Jeff Without waiting for a human to click the button.

Emily Right. And finally, the fifth pillar is learn. The system continuously ingests the outcomes of all these investigations, whether they were automated or human driven, folding that telemetry back into its foundational model. So every single future decision is sharper and more precise.

Jeff Assist, investigate, build, act and learn. It is just a complete self-reinforcing life cycle of defense.

Emily It really.

Jeff Is. The future of financial security isn't about exhausted human analysts translating their hunches into endless manual clicks and reactive rule writing. It is about powerful AI agents seamlessly translating human intent into immediate, platform wide execution.

Emily It equips the defenders with the speed and automation that the syndicates have been exploiting for years, which.

Jeff Is a genuinely thrilling vision for the future of global finance. It feels like the infrastructure is finally catching up to the press. Thank you for joining us on this deep dive. We'll see 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.