EPISODE 27 | Monthly Brief

Fraud in the AI Era

July 13, 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.

Emily Right now. Like as you listen to this, an entirely fabricated AI generated human being is probably applying for a mortgage at your bank.

Jeff It's a pretty terrifying reality. Yeah.

Emily And I mean, the really crazy part, the bank's traditional security systems are fundamentally designed to just, you know, hold the door open for them, right?

Jeff Because the contrast is just jarring when you think about it. I mean, the cultural idea of a bank heist involves ski masks, right? A getaway car.

Emily Sirens ringing down the street.

Jeff Exactly. You can see it. You can quite a camera at it. And you can design physical security to stop it. But the reality of modern financial crime is completely silent. It's happening entirely in the servers that power our digital lives.

Emily We are talking about a hidden, incredibly high stakes cat and mouse game of global money movement. Literally billions of dollars on the line. And for the financial institutions guarding that money, well, the rules of engagement just dramatically changed. So welcome to the deep dive.

Jeff Thanks. Glad to be here for this one.

Emily Yeah. We've got our hands on some truly fascinating intelligence today for you listening. It's a dispatch from a July twenty twenty six DataVisor industry newsletter, and it is detailing a massive shift in how global money movement is policed. We're looking at huge regulatory crackdowns, multi-million dollar fines, and this emerging AI arms race happening right behind the scenes of your bank account.

Jeff Yeah. And the core mission we really need to tackle today is understanding that friction. We need to look at why the old ways of doing things, what the industry calls disconnected fraud and anti-money laundering or AML systems, why those systems build for a different era are suddenly becoming a massive liability.

Emily Like an existential threat to these banks.

Jeff Totally. This isn't just about banks doing more paperwork to keep regulators happy anymore. It is a fundamental shift in how we define track and intercept suspicious human behavior, and an environment where, honestly, the human might not even be real.

Emily So let's unpack the consequences of failing to adapt to this new environment, because the intelligence points to a glaring multi-million dollar example of what happens when an institution gets it wrong.

Jeff Right? The Merrill Lynch case.

Emily Exactly. The SEC recently slapped Merrill Lynch with a seven point five million dollars fine. And the reason given is that they failed to file hundreds of suspicious activity reports, or SARS. Hundreds of millions of dollars just slipped right through the cracks of their system.

Jeff And the mechanism of that failure, I think, is what tells a real story here. The SEC found that Merrill Lynch was relying on automated transaction monitoring systems that were built on rigid risk thresholds.

Emily Meaning what exactly?

Jeff Well, they were using static rules to govern highly dynamic environments.

Emily Let me make sure I have this conceptualized. Right. Because static rules based transaction monitoring sounds, uh, incredibly clinical. It feels a bit like a nightclub bouncer who is strictly given a rule to, I don't know, only stop people wearing sneakers.

Jeff That's actually a great way to think about it.

Emily So the bouncer stands there all night just staring at everyone's feet. Meanwhile, a guy in dress shoes walks right through the front door carrying a crowbar, and the bouncer just waves him in because, hey, he's not wearing sneakers. The rule was followed perfectly, but the actual threat was completely missed.

Jeff Exactly, because rigid thresholds are entirely blind to context. They look for very specific, pre-programmed scenarios, like if a bank sets a rule to flag any transaction over ten thousand dollars, the system only looks for ten thousand dollars, right? But financial crime grows more sophisticated by the day. Fraudsters know where those static thresholds are, so they intentionally structure their transactions, say, moving nine thousand nine hundred and ninety nine dollars to fly just under the radar of that specific hard coded rule.

Emily I have to play devil's advocate here for you listening, though, because looking at a seven point five million dollars fine, it's easy to wonder, isn't this just regulators trying to squeeze fines out of big banks for minor technicalities? Like, is it really that the banks are negligent, or are the criminals just getting way too fast for human compliance teams to catch.

Jeff Well, what's fascinating here is that the regulators are actually responding to that exact reality. They know the criminals are faster. Which brings us to a really pivotal moment outlined in our sources. July eight, twenty twenty six.

Emily It's a big overhaul.

Jeff All right. The Federal Reserve officially joined FinCEN, the FDIC and the OCC in proposing amendments to their AML program rules. It's widely considered the most significant overhaul of the Bank Secrecy Act in over twenty years.

Emily Wow. Twenty years.

Jeff Yeah. And the regulators are essentially outlawing that lazy bouncer approach. They are stating that relying on static rules when you know criminals are evading them is in fact a form of negligence.

Emily Okay. So if the old way is officially dead, what are the feds actually demanding now? The source points to three massive shifts coming from this Bank Secrecy Act overhaul. The first one is explicitly integrating countering the financing of terrorism or CFT into their frameworks. Right. And this isn't just a secondary compliance check anymore. It has to be treated as a core legal requirement.

Jeff Yeah. And integrating CFT changes the fundamental architecture of a compliance program. Yeah. Historically, money laundering and terrorism financing were treated with slightly different urgencies.

Emily How so?

Jeff Well, money laundering is usually about taking dirty money and making it look clean. Terrorism financing is often the complete opposite.

Emily Taking clean money and making it dirty.

Jeff Basically, yeah. Taking clean money maybe from legitimate charities or businesses and hiding its destination to fund illicit acts by making CFT explicit. Regulators are mandating that institutions actively hunt for both directional flows simultaneously as a primary objective.

Emily So if counterterrorism is now the baseline, that implies banks can't just set up a defense once and forget about it. How are they supposed to actively hunt these moving targets without going bankrupt from compliance costs?

Jeff That's the million dollar question.

Emily Which seems to connect to the second big change, right? This shift to a dynamic risk based allocation of resources. The feds are moving away from rigid. Check the box. Compliance. Financial institutions have to continuously evaluate their specific threat landscape.

Jeff Exactly. Because checking the box used to mean running a standardized report every thirty days, filing the necessary paperwork, and just calling it a day. Right. Dynamic resource allocation is entirely different. It requires a bank to shift its investigative power on a Tuesday, simply because a new threat vector appeared on Monday. You are legally required to proactively hunt for risks and move your tech to your highest risk areas in real time. You can't just wait for an outdated threshold to trip an alert.

Emily And the third major takeaway from this July eighth overhaul is that the clock is ticking. The fed is now fully aligned with FinCEN and the other agencies presenting this unified regulatory front.

Jeff Yeah, standard operating procedures have to be updated right now.

Emily Right ahead of the final rulemaking. There's no grace period for relying on those old siloed systems.

Jeff The urgency is definitely palpable. Institutions just cannot hide behind fragmented internal departments anymore. If one hand doesn't know what the other is doing, the regulator will hold the entire entity accountable.

Emily Which brings us to a massive structural problem for the banks. If they can't use the old static systems and they have to constantly adapt, what is the actual cure? I mean, if you fire the bouncer staring at the sneakers, who do you replace him with? Right? The intelligence points heavily toward tearing down the internal silos, specifically moving toward something called FRAML.

Jeff Yes, FRAML. It stands for the unification of fraud and anti-money laundering operations.

Emily Which used to be totally separate.

Jeff Completely isolated. Historically, these two departments within a bank operated like two different companies under the same roof. The fraud team focuses on money being stolen from the bank or its customers, like someone hacking an account or using a stolen credit card.

Emily In the AML team.

Jeff They look for dirty money being moved through the bank by criminals because they have different goals. They historically use completely different systems, different data sets, and entirely different workflows.

Emily But the criminals don't care about the bank's internal org chart. I mean, the person committing the fraud to steal the money is often the exact same person, or at least part of the same network. Laundering that stolen money.

Jeff Exactly. And having siloed systems creates massive operational inefficiencies and really dangerous blind spots. When data is trapped in separate systems, you get a hugely high rate of false positives. Investigations take weeks because analysts have to manually gather data from multiple disjointed screens just to piece together a single customer's behavior.

Emily It sounds like a nightmare, but we actually have a concrete case study here to see how this unification works in practice. NASA Federal Credit Union.

Jeff Oh yeah, they just won the twenty twenty six Celent Model Risk Manager of the year award.

Emily Right. And before this, they were suffering from exactly what you just described siloed systems creating high false positives, blinding them to evolving crimes, bogging down their investigations. And they solve this by deploying Data Visor's Unified framework platform.

Jeff Yeah. By consolidating fraud detection, AML monitoring, case management, and real time analytics all into a single platform powered by AI. They eliminated the blind spots. The fraud team and the AML team are finally looking at the same holistic picture of customer behavior.

Emily The results they achieved are honestly staggering. But I want to break down what these numbers actually mean for the real people involved. By moving to this unified AI platform, NASA federal credit unions saw a forty two percent reduction in AML. False positives.

Jeff That's massive.

Emily They also saw a forty one percent decrease in manual review time, a twenty percent improvement in fraud detection, precision and their suspicious activity, reports the SARS that Merrill Lynch got fined for missing. They are now being filed ninety percent faster.

Jeff To to put that forty two percent drop in false positives into perspective for you listening every false positive is a real human being whose bank account was suddenly frozen while they were, you know, trying to buy groceries or pay their rent.

Emily Oh yeah. I hadn't thought about it like that.

Jeff A false positive isn't just a metric on a dashboard. It is immense friction for a legitimate customer. So NASA FCU essentially saved thousands of their members from that humiliation and frustration.

Emily Well, the newsletter also mentions they achieved a fifty to sixty percent reduction in operating and platform costs, which equates to about two hundred thousand dollars in annual compliance savings. Doug Nahhas, the COO of NASA FCU, noted that this let them focus on serving their members instead of just managing risk.

Jeff Which is the ultimate goal, really.

Emily But I have to push back a little here.

Emily Saving two hundred thousand dollars in operating costs is, you know, a nice bonus. But considering Merrill Lynch just got fined seven point five million dollars, isn't the true return on investment of a financial platform simply staying out of the SEC's crosshairs?

Jeff If we connect this to the bigger picture? Yeah, avoiding a multimillion dollar fine and the reputational damage is certainly the ultimate ROI. However, the true technical magic of the NASA FCU story lies in a different metric entirely. They achieved that forty two percent drop in false positives, while retaining one hundred percent of their true positive weight.

Emily Really, keeping one hundred percent of the true positives means they didn't miss a single actual crime while ignoring almost half of the false alarms. That sounds almost impossible. How does the system know what to ignore without accidentally letting a threat through the mechanism?

Jeff Driving that accuracy is unsupervised machine learning, or UML unsupervised? Yeah, unlike our bouncer looking for specific sneakers or a static rule looking for a ten thousand dollars transfer, unsupervised machine learning doesn't need to be told what a crime looks like before it can catch it. Okay? It doesn't rely on labeled historical data. Instead, it looks at the invisible webs connecting vast amounts of data in real time.

Emily So how does that actually work in practice? Like if it's unsupervised, what is it looking for?

Jeff It clusters behaviors to find outliers. For example, it might see that a new account was opened by someone claiming to live in Ohio, but the UML algorithms detect that the velocity of the mouse movements on

the web page. The device fingerprints and the IP routing all match the behavioral patterns of a known fraud ring operating out of Eastern Europe.

Emily That is wild, right?

Jeff It flags the subtle, complex relationships between data points, how fast funds are moving, where they are going, the specific devices being used, and it identifies anomalies that a human would frankly never catch.

Emily And it does this without needing new rules written.

Jeff Exactly. It adapts to the threat landscape dynamically. Right. Completely satisfying the new Federal Reserve regulations without requiring human engineers to constantly write new rules.

Emily So the technology is incredibly powerful. And obviously, NASA f.cuz success proves it works. But the intelligence we're reviewing makes it very clear that this level of power isn't just a luxury for the most tech forward banks anymore. We are officially entering an era of AI versus AI warfare.

Jeff And it's playing out on a global stage, too.

Emily Yeah, the regulatory pressure is definitely international. In Canada. The office of the Superintendent of Financial Institutions, Osfi, has rolled out specific twenty twenty six expectations requiring Canadian institutions to handle dual regulator scrutiny.

Jeff What is particularly vital in those Canadian guidelines is their intense focus on model explainability and vendor governance.

Emily Let's look at the attacker side first, just to understand why regulators like OSFI are pushing so hard for advanced defenses. Because according to a twenty twenty six generative AI webinar highlighted in our sources, fraudsters have evolved way past sending, you know, badly spelled phishing emails.

Jeff Way past.

Emily That. They are using gen AI to completely automate scams and bypass KYC. Know your customer controls.

Jeff The scale of the attacks is fundamentally changed. We are dealing with industrial scale fraud powered by synthetic identities.

Emily I've heard that term a lot. Synthetic identities. What does that actually mean? Are they just stealing someone's driver's license?

Jeff No. It is far more sophisticated than traditional identity theft. Generative AI allows a fraudster to create what's basically a Frankenstein identity.

Emily A Frankenstein identity like stitching parts together.

Jeff Exactly. They might take a real Social Security number, but pair it with a fabricated name, an AI generated face, and a completely synthetic credit history. And because parts of the identity are real, it can easily bypass basic checks. A single fraudster using gen AI can spin up thousands of these entirely fabricated human beings, complete with deep faked video verification documents, and just launch them at a bank's onboarding system simultaneously.

Emily So a human just gets completely overwhelmed.

Jeff Bright human investigators simply cannot click through those alerts fast enough to manually verify them.

Emily It's less like fighting a gang of robbers and more like dealing with a biological immune system under attack. The generative AI attacks are like a rapidly mutating virus, just overwhelming the body by sheer volume and adaptation.

Jeff Which brings us to the defenders side. Right? The immune system creating custom antibodies in real time. The DataVisor dispatch highlights an upcoming July twenty three webinar focusing on the rise of Agentic AI in financial institutions.

Emily Hold on Agentic AI that sounds heavily like a Silicon Valley buzzword. What makes this different from the standard chatbots banks have been using for years? That, let's be honest, usually just frustrate customers trying to find their routing number.

Jeff I know it sounds like jargon, but a standard chatbot is reactive, right? It waits for a prompt and provides a pre-programmed response. An Agentic AI acts as an autonomous agent.

Emily Okay, so it does things on its own.

Jeff Yes. Financial institutions are deploying these systems to automate rule optimization, streamline complex workflows, and eliminate manual triage entirely. The defensive AI is actively analyzing the incoming attacks from the fraudulent AI and adapting the institution's defenses on the fly, without waiting for human permission to update a static rule.

Emily Here's where it gets really interesting, though. I have a major practical question about this. Thinking about you as a listener who actually has to implement or audit these systems. Yeah. If an institution is using an autonomous AI that optimizes its own rules on the fly, constantly learning and shifting its parameters to catch these mutating genii viruses. How on earth do they pass that Canadian OSFI requirement for model explainability?

Jeff That is the question, right?

Emily How do you explain an AI's decision to a human auditor who wants a nice, neat linear flowchart? If the AI is constantly teaching itself in ways we barely even understand.

Jeff You have just hit on the central tension of modern fin crime. It is the razor's edge. Every single institution is walking right now. How do you deploy autonomous defense without surrendering accountability? Right? And this is exactly why those twenty twenty six OSFI guidelines heavily emphasize transparent decision making and unified data architecture.

Emily Because you can't just tell a regulator, well, the mysterious black box said the transaction was bad, so we blocked it. But we have no idea why exactly.

Jeff A black box is a massive regulatory liability. The AI must be highly advanced to catch the synthetic identities, but it simultaneously has to show its work.

Emily Show its work. Like in math class.

Jeff Basically, yeah. Modern platforms are designed so that when an AI flags a complex synthetic identity ring, it automatically generates a transparent, human readable narrative of its logic. It maps out the behavioral intelligence, the real time patterns, the device fingerprints, and presents it clearly in plain language.

Emily So it actually writes the report for the human.

Jeff Precisely. The human investigators need that transparent logic so they can actually understand the threat, confidently write their own assessments, and file those suspicious activity reports. The AI is doing the heavy lifting of detection and triage at machine speed, but the system architecture must be designed to empower human accountability, not replace it.

Emily So let's wrap this up. What does this all mean for you listening? We started this journey by looking at the collapse of the old ways, the reliance on static rules and rigid thresholds, the lazy bouncer staring at the sneakers essentially. Right. And that is literally resulting in multi-million dollar SEC fines like we saw with Merrill Lynch, because it simply cannot see the modern threats.

Jeff Blind.

Emily To them. Exactly. The regulators from the Federal Reserve overhauling the Bank Secrecy Act to OSFI in Canada, they are legally demanding dynamic, real time risk management that hunts for both money laundering and terrorism financing. And we explored how unified AI platforms, tearing down those internal silos with raml are answering that call.

Jeff And providing massive efficiency gains to.

Emily Yeah, saving institutions like NASA, Federal Credit Union, hundreds of thousands of dollars while protecting innocent customers from having their accounts wrongfully frozen. And finally, we peeled back the curtain on the terrifying but fascinating arms race currently underway between generative AI fraudsters launching synthetic identities and agentic AI defenders, mutating the bank's rules in real time to stop them.

Jeff The landscape is just shifting beneath our feet so fast. I want to leave you with one final thought to mull over. Building on the trajectory of everything we've just unpacked from this intelligence.

Emily Lay it on us.

Jeff We know fraudsters are using gen AI to bypass identity controls at an unprecedented scale. And we know institutions are deploying autonomous AI to dynamically catch them, translating complex behavioral webs into human readable alerts. But what happens when the fraudulent AI and the defense of AI start learning directly from each other? Are we fast approaching a future where the actual battle of financial crime is fought entirely in milliseconds, machine to machine, with human investigators simply arriving at their desks the next morning to read the AI generated after action reports of a war that was won or lost while they slept.

Emily Oh wow. It is a wild thought and honestly, a reality that might be closer than we think. As these systems continue to evolve. Thank you for joining us on this deep dive. Keep questioning the systems. Working silently behind the scenes of your everyday world. The hidden cat and mouse game of global finance is still happening. The mice just learned how to code and the cats are upgrading their hardware. We'll 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.