EPISODE 19 | Analysis

Fraud & AML Case Management Playbook for 2026

MAY 18, 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 Imagine for a second that your job is to find a needle in a haystack. But there's a catch, right?

Jeff Someone is constantly dumping more hay on you.

Emily Exactly. And the hay itself is like constantly changing shape. Plus, if you happen to miss the needle, well, the government might find your company millions of dollars.

Jeff I mean, it sounds like an impossible stress dream, honestly. But for financial crime investigators working in 2026, that's literally just a Tuesday morning.

Emily Yeah. And today we are doing a deep dive into why that system feels so fundamentally broken. And, you know, more importantly, how high performing teams are actually fixing it.

Jeff I've got some really fascinating source material for this.

Emily We do. We are unpacking this comprehensive playbook by Fang Yu from DataVisor. It's basically a roadmap detailing why fraud and anti-money laundering case management systems are just they're breaking down under modern pressure.

Jeff And it outlines this seven step framework to actually survive it.

Emily Right. So our mission today is to cut through all the industry jargon, figure out how financial crime has evolved and understand the mechanics of why the old tools just, you know, cannot keep up anymore.

Jeff To do that effectively, though, we really need to establish a quick baseline because fraud and anti-money laundering or AML, they're two very different functions.

Emily Even though they operate in the same ecosystem.

Jeff Exactly. I mean, fraud asks this very operational immediate question, right? Like, did this person steal something? And how do we stop the financial loss right now? It is completely about stopping the bleeding.

Emily Like, uh, freezing an account or blocking a wire transfer. The literal second, it looks suspicious.

Jeff Yes, but AML, on the other hand, asks a regulatory historical question. It's asking, is this activity indicative of money laundering and do we need to report it to FinCEN?

Emily So it's more about mapping out the broader criminal network to meet strict government compliance.

Jeff Right. Two completely different mandates. But as the DataVisor playbook shows, they are colliding in ways that legacy tech simply wasn't built to handle.

Emily You know, reading through this playbook and looking at all these fractured legacy systems, it actually really reminds me of Kintsugi.

Jeff Oh, the Japanese pottery thing.

Emily Yeah, exactly. That art practice where if a ceramic bowl shatters, they don't just throw it away and they don't glue it back together to try and hide the cracks.

Jeff Either they rejoin the pieces with gold, right?

Emily Right. Gold or silver lacquer. They actively highlight the fracture points and fill them with something stronger. So the repaired bowl is actually more resilient than the original ever was.

Jeff That is a phenomenal way to frame it.

Emily That's totally what this playbook feels like. We aren't taping a broken case management system back together. We're identifying the specific points of failure and filling them with entirely new workflows.

Jeff Yeah. And to understand how to apply that gold lacquer, we first have to understand why the bowl shattered in the first place.

Emily Which brings us to the first major structural crack, right?

Jeff Investigators are often failing before they even open a case.

Emily Because they are completely drowning in noise. I mean, the source notes that legacy systems are generating false positive rates sitting anywhere between seventy percent and ninety five percent.

Jeff It's staggering.

Emily If you're listening to this and you work in any kind of triage it. Customer support medicine. Just imagine if ninety five out of every one hundred alerts you got were complete waste of time.

Jeff You would develop alert blindness instantly. You just start clicking ignore.

Emily Exactly.

Jeff And that alert overload, it's actually a symptom of a much deeper problem. It's this shift in how financial crime actually operates.

Emily Now how so?

Jeff Well, the old playbook was built for transaction monitoring. The system looked for a single bad actor doing a single bad thing, like a massive transfer over ten thousand dollars, or a login from a foreign country at three o'clock in the morning.

Emily It was very rules based, very binary.

Jeff Exactly. But criminals know those rules now.

Emily They read the same playbooks.

Jeff They do, and they adapted. Financial crime in twenty twenty six isn't based on single loud events anymore. It is based on coordinated, quiet networks.

Emily Are like synthetic identity rings.

Jeff Yes, fraudsters will invent a person out of thin air, get a credit card, and spend months using a completely. Normally, they pay the bill on time. They buy groceries.

Emily Just to establish a baseline of quote unquote normal behavior.

Jeff Right. And then they use massive mural networks to move money in these tiny structured increments that fly completely under the radar of those old binary rules.

Emily So catching them requires a total shift to behavioral surveillance.

Jeff You have to understand what normal looks like for an entity over a long period of time, and flag the really subtle deviations.

Emily Which completely explains why reviewing an individual transaction alert is so useless. Now. It feels like looking at a single piece of a jigsaw puzzle through a microscope and trying to guess the whole picture.

Jeff It really does. I mean, if a sophisticated fraud ring is using forty different accounts and fifteen different devices across three states, a human being is never going to connect those dots when they are buried under ninety five percent false alarms.

Emily The human brain simply cannot process disparate data points at that scale or speed, for that matter.

Jeff No, the system is asking analysts to do work that should never reach a human desk. So the playbook proposes a two part structural fix here intelligent alert triage and entity level investigation.

Emily Okay, break this down for me.

Jeff First, you place an AI pre-filtering layer in front of the human cue. The AI analyzes the alerts against historical behavioral patterns and basically auto closes the low risk noise before an analyst ever even sees it.

Emily Okay, let me pause you there because my immediate thought is regulatory blowback.

Jeff Oh for sure.

Emily Right. If an AI is just like silently closing security alerts in the background, regulators are going to have a field day with that lack of oversight.

Jeff They absolutely would, which is why the mechanism requires the AI to generate mandatory reason codes for every single auto closure.

Emily Oh, so it's not silent.

Jeff No it doesn't just delete the alert. It appends a log saying closed because device ID has a three year history of legitimate travel to this IP address or whatever the specific reason is, it is fully documented.

Emily Okay, that makes sense. But what about the alerts that don't get auto closed?

Jeff The real magic happens for the alerts that do reach the human, the five percent that are actually highly suspicious, right?

Emily When the analyst opens that high risk case, what do they see?

Jeff They aren't looking through your microscope anymore. The system uses graph visualization instead of a spreadsheet of individual transactions. The analyst sees a literal web on their screen.

Emily A web like a spider web.

Jeff Sort of a central node might be an IP address, and radiating out from that node are lines connecting it to every user account, every device fingerprint, and every physical address that has ever interacted with that IP.

Emily Oh, wow. So instead of playing whack a mole with forty different isolated alerts, the investigator sees the entire fraud ring in one single view.

Jeff Yes, they can highlight the entire web and shut down all forty accounts simultaneously.

Emily The source material actually claims. This entity level approach can create efficiency gains of up to 100x.

Jeff Which sounds like marketing fluff until you actually see it.

Emily Right? That is a staggering number. But when you explain the mechanics of visualizing the whole network at once, it actually makes total sense.

Jeff It changes the fundamental nature of the investigation from reactive to proactive. But, um, once you zoom out and give investigators the ability to see that full entity, a glaring new problem pops up.

Emily Which.

Jeff Is you realize the right hand has no idea what the left hand is doing. The fraud team and the AML team are completely siloed.

Emily Ah, the classic corporate silo. But I imagine in this context, it isn't just inefficient. It's actively dangerous.

Jeff It is. The playbook points out that fraud is very often the predicate offense to money laundering, right?

Emily Because the money has to originate somewhere. So the proceeds of like an authorised push payment scam and app scam, they don't just vanish into thin air.

Jeff No, the criminals immediately move those stolen funds into their mule networks to be laundered.

Emily And yet these teams are often working in completely fragmented environments. I think the playbook notes institutions are frequently using four to six separate tech stacks across risk departments.

Jeff Which means a fraud investigator might catch an account takeover, freeze the account, close their case and celebrate a win.

Emily Meanwhile, they have absolutely no idea that the AML team down the hall desperately needs that exact, compromised account data to track where the rest of the laundered money is flowing.

Jeff Exactly. The criminal just slips right through the gap between the two departments.

Emily Okay, look, if the criminal is explicitly exploiting the gap between fraud and AML, why are we tiptoeing around this? Why not just smash the two teams together?

Jeff I know it seems obvious, right?

Emily Put them in the same room under one manager with one budget, one giant financial crime, super squad. Problem solved. Right.

Jeff It is incredibly tempting to view merging as the silver bullet, but you absolutely cannot do it. The regulatory frameworks forbid it for a very practical reasons.

Emily Why?

Jeff Remember our baseline definitions from the intro.

Emily Operational versus regulatory.

Jeff Right. The fraud team is operating under intense pressure to stop immediate operational loss. They have to freeze funds in milliseconds. The AML team is operating under strict government mandates to map networks and file suspicious activity reports to FinCEN within a rigid thirty day window.

Emily Oh I see.

Jeff If you merge the teams, you bleed those mandates together. An analyst focused on stopping a five hundred dollars fraud loss might inadvertently tip off a massive money laundering suspect.

Emily Which could ruin a federal investigation.

Jeff Exactly. They have fundamentally different thresholds for action.

Emily So you are trapped. You can't merge the teams. But keeping them siloed lets the criminals win. How do you share the data without breaking compliance?

Jeff You implement what the industry calls frames the convergence of fraud and AML. But you have to do it correctly.

Emily And what does correctly look like here.

Jeff The architecture requires separate workflows, but a shared unified intelligence layer. The teams do not work each other's cases. However, the data flows natively between them in the background.

Emily Okay, so how does that work in practice?

Jeff When an AML investigator opens a profile to investigate potential laundering? The platform automatically queries the shared intelligence layer and surfaces the fraud team's historical signals on that exact same entity.

Emily Oh, so there's no copy pasting between systems, no sending emails to ask for case files.

Jeff Exactly. They both feed into and draw from the exact same data well, but they view that data through their own specific regulatory lenses.

Emily Okay, so we fixed the alert noise and we have the teams sharing data through a unified layer. We are building a really solid foundation.

Emily Here we are.

Emily But let's say this unified team catches a massive new app scam. The fraudsters are watching. They realize their vector is burned. So they pivot to a completely new tactic overnight.

Emily Which.

Jeff They always.

Emily Do, right?

Emily If the underlying software system is rigid, does all this shared data even matter if we can't update our detection rules?

Jeff This is breakdown number three. The adaptation gap. And it is driven by what practitioners call the engineering bottleneck.

Emily Oh man.

Emily If you are listening to this and you have ever had to submit an IT ticket just to change a basic drop down menu or update a parameter in your CRM, you know exactly the pain we are about to describe.

Emily It is that exact pain.

Jeff But with millions of dollars on the line in legacy hard coded systems, risk analysts will spot a brand new attack vector in real time, but they do not have the technical access to adjust the detection rules themselves. They have to file an engineering ticket.

Emily And that.

Emily Ticket goes into a software development sprint cycle.

Emily Right?

Jeff The new code has to be written, placed into a staging environment, run through quality assurance testing.

Emily And suddenly weeks or even business quarters have gone by. It's like realizing the lock on your front door is broken, but instead of just going to the hardware store and fixing it, you have to submit a permit to the city. Wait for an inspector and the approval takes three months.

Emily That's exactly what it's like.

Emily Meanwhile, the fraudsters are just casually walking in and out of your house, exploiting the known vulnerability.

Emily Yep.

Emily So how do institutions close this adaptation gap without swinging the pendulum too far? Because giving Non-engineers the power to change live algorithmic rules sounds like a great way to accidentally flag every single transaction in the bank and shut down the business.

Jeff It is a massive operational risk. The playbook solves this through a combination of no code rules and champion challenger model learning.

Emily Okay, what are no code rules?

Jeff First, the platform provides a graphical interface where risk teams can configure and update rules themselves using drop down logic, rather than writing Python or Java. They can respond to a threat in hours.

Emily But to prevent them from accidentally breaking the bank.

Emily The.

Jeff System enforces shadow testing.

Emily Okay, walk me through the mechanics of shadow testing. How does it actually protect the live environment.

Jeff When an analyst creates a new rule, let's call it the challenger rule, it does not go live. The system runs it in the background, in the shadows, against a massive data set of historical transactions.

Emily Oh, okay.

Jeff Simultaneously, it runs the existing champion rule against that exact same data. The system then compares the deltas.

Emily So it generates a mathematical report.

Emily Exactly.

Jeff It shows the analyst. If this challenger rule had been live for the last ninety days, it would have caught fifteen percent more fraud and the false positive rate would have dropped by two percent.

Emily Oh I see, so you proved the math on historical data before it ever touches a live customer account.

Emily Precisely.

Jeff Only when the challenger mathematically defeats the champion does it get promoted to the live environment.

Emily That's incredibly smart.

Jeff And to ensure the machine learning models underlying all of this continue to get smarter, the playbook emphasizes sub typology tagging.

Emily Sub typology tagging. What does that actually mean in practice?

Emily Well, in a bad system.

Jeff When an analyst closes a case, they select a generic tag from a drop down menu like account takeover.

Emily Right.

Jeff The AI just learns that something really bad happened, but in a high performing system, the tags are hyper granular.

Emily Like what?

Jeff The analyst tags it as account takeover via SIM swap, or account takeover via credential stuffing. That granularity feeds back into the continuous learning model, telling the AI exactly which specific behavioral features to weight more heavily the next time it scans the network.

Emily So.

Emily It is constantly refining its own parameters based on the specific outcomes of closed cases. But taking a step back here, we are talking about letting risk teams deploy models in real time and using algorithms that learn dynamically. If we are giving the system this much autonomy, won't regulators eventually freak out about a black box making these calls?

Emily That is arguably.

Jeff The most urgent issue in the entire playbook. It's driven by the FinCEN April twenty twenty six Notice of Proposed rulemaking.

Emily Are the twenty twenty six reckoning. What exactly is changing in the regulatory landscape?

Jeff Regulators are fundamentally shifting how they evaluate compliance programs. For years, it was check the box compliance. You proved to the examiners that you purchased a sophisticated transaction monitoring system. And that was mostly enough.

Emily Basically just showing you had.

Emily The tool, right.

Jeff But the twenty twenty six rules demand demonstrated effectiveness. Examiners do not just want to know that your AI flagged an entity. They want to know the exact mathematical and behavioral reasoning behind the flag.

Emily Okay, let me play devil's advocate here. We implement AI specifically because it can see multidimensional patterns across millions of data points that the human brain literally cannot comprehend. Right. So if an advanced neural network flags a high value customer with a massive risk score of, say, ninety eight out of one hundred, why are we handicapping the AI by demanding it only use logic a human can understand? Shouldn't the analyst just trust the math and file the report?

Jeff Trusting the math without understanding it results in what the industry calls the Undefendable decision.

Emily The.

Emily Undefendable.

Emily Decision.

Jeff Yeah, if you have a risk score of ninety eight, but the system is a black box and cannot show you the reasoning path, it is a massive regulatory liability.

Emily Because you have to defend it to an auditor.

Emily Exactly.

Jeff Imagine your analyst is sitting across from a federal examiner. The examiner asks, why did you file a suspicious activity report on this specific customer? If the analysts only answer is, uh, the computer gave them a ninety eight. That is an immediate program deficiency.

Emily Yeah, that makes sense. You cannot audit an algorithm you don't understand. So how does the DataVisor framework actually solve the black box.

Emily Through.

Jeff Explainable AI? The system is architected to produce a human readable reasoning path for every single decision.

Emily Like.

Emily A plain English summary.

Emily Yes.

Jeff It translates the algorithmic weights into a clear narrative, like we flag this entity because their device fingerprint matches a known neural network. Their transaction volume spiked four hundred percent on a Tuesday, and they bypassed a geolocation check.

Emily Wow. That's specific.

Jeff It has to be. Yeah. And crucially, every one of those claims must be hyperlinked and traceable back to the raw underlying transaction logs. You have to show your work.

Emily That makes total sense. You have to be able to defend the logic. But the playbook mentions another trap here alongside the black box. The audit gap.

Jeff Ah yes, the audit gap is basically a problem of time.

Emily How so?

Jeff Let's say an auditor asks to see why an analyst made a specific decision on an account six months ago. In a legacy system, the compliance officer pulls up the account, but the dashboard shows the account as it exists today.

Emily Wait, really?

Jeff Yeah. The customer has updated their address. New transactions have posted. Old alerts have cleared. The context the analyst saw six months ago is completely gone.

Emily Oh my gosh. You cannot prove they made the right call because you can no longer see what they saw. That is a terrifying vulnerability. You could make the perfect decision based on the data you had on a Tuesday, but if the data changes by Friday, you look incompetent to an auditor.

Jeff To fix this, the system must capture an immutable audit trail snapshot.

Emily How does that work?

Jeff When the analyst clicks, approve or close on a case, the platform literally takes a digital picture of the exact data, state, the specific alerts that were visible, and the precise risk scores at that exact second in time.

Emily And it just saves it.

Jeff It locks that snapshot in a vault. So when the examiner asks about it six months later, you can instantly recreate the exact visual environment the analyst was operating in.

Emily That alone would save compliance officers countless sleepless nights.

Jeff Without a.

Emily Doubt. Which actually brings us to the final hurdle. We have investigated the entity. The A's reasoning is fully explained, and we have an immutable snapshot locked in the vault. Right? But if the actual compliance paperwork isn't filed correctly and on time, this entire sophisticated process was basically for nothing.

Jeff And this is the seventh and final breakdown. Treating compliance reporting as an administrative afterthought.

Emily I can completely visualize this workflow breaking down like an investigator spends a week untangling a massive web of synthetic identities, they finally close the complex case, and then they open a blank word document and stare at a blinking cursor. Yep.

Jeff Trying to remember all the intricate details from memory to write up a suspicious activity report or, you know, a written statement of unauthorised debit at w s u d for a fraudulent ACH return.

Emily And human memory is incredibly fallible.

Jeff Highly fallible. Reconstructing complex network cases days later leads to thin narratives, inconsistent details, and, most dangerously, late filings.

Emily Because FinCEN enforces a strict thirty day filing deadline for SARS.

Jeff Right, exactly. If you treat the report as a postscript, you will miss deadlines. And under the twenty twenty six Demonstrated effectiveness rules, late filings and poor quality narratives will result in severe examiners findings and potential fines.

Emily So how do you fix it without just hiring an army of technical writers to follow the analysts around?

Jeff You integrate the output, you build the compliance report during the investigation itself.

Emily Oh, like simultaneously?

Jeff Yes. The system should automatically assemble the SAR narrative or the regulatory documentation as a structured byproduct of the investigator's natural workflow.

Emily How seamless is that, actually?

Jeff It's very seamless. Every time the analyst highlights a transaction in the graph and tags it as anomalous, the system auto populates that data into a chronological narrative field in the draft SAR.

Emily Oh, wow.

Jeff And when they link two entities together visually, the system automatically translates that relationship into text for the report. By the time the investigation is done, the report is already ninety percent written.

Emily That is incredibly efficient. But is it fully automated?

Jeff No. And there is a non-negotiable rule here. There must be a human in the loop for final approval.

Emily Because regulators won't accept an AI doing the filing.

Jeff Exactly. Regulators will not accept autonomous AI submitted SARS. The AI acts as the drafter pulling together the evidence, but the human analyst acts as the editor in chief. They review the narrative, ensure the explainable AI logic makes sense, and they are the ones who officially file it.

Emily Okay, we've covered an immense amount of ground here today. We've outlined seven massive structural breakdowns, from drowning in a loud noise to the siloed departments, the engineering ticket trap, black box algorithms, and botched reporting.

Jeff There's a lot to take in.

Emily It really is. If someone is listening to this right now, sitting at their desk and feeling totally overwhelmed by the sheer scale of the problem. Where on earth do they start applying the gold lacquer to their kintsugi bowl?

Jeff The DataVisor playbook is very clear on triage. You sequence your investment based on your highest operational pain, not by whatever vendor future looks the most appealing. You must triage the highest bleed.

Emily Meaning don't just buy the shiny new graph visualization tool because it looks cool in a demo.

Jeff Exactly. I mean, if your analysts are burning out and your queue never empties, your highest bleed is alert overload. Do not buy a flashy graph tool first, because your team won't even have the time to look at the graphs anyway, right?

Emily That makes total sense.

Jeff You need to invest in intelligent alert, triage, and AI pre-filtering to stabilize the queue first.

Emily What about reporting issues?

Jeff Conversely, if your SRS are constantly being returned by quality assurance for poor narrative quality, your bleed is reporting. Start with integrated compliance outputs. You have to fix the specific foundation that is currently sinking your operation.

Emily That is intensely practical advice. Well, let's bring this all together. Today we journeyed from the nightmare of surviving ninety five percent false positive alert queues to the necessity of building a proactive, explainable, and unified friemel ecosystem.

Jeff A huge leap.

Emily Absolutely. We've explored the mechanics of why old transactional rules cannot catch coordinated entity level networks, and how graph visualization exposes them.

Jeff And we learned that the twenty twenty six regulatory standard is entirely about demonstrated effectiveness.

Emily Right? You cannot just possess the tools anymore, you have to prove mathematically and visually exactly why your system works using immutable audit trails and human readable AI.

Jeff It is a massive paradigm shift for the industry. You are moving from a posture that is inherently reactive, siloed, and opaque to an architecture that is continuous, converged and fully defensible to an examiner.

Emily And as we wrap up this deep dive, I want to leave you, the listener, with a final thought to ponder. We've established today that regulators now demand highly transparent, explainable AI. The system has to clearly and explicitly document in human readable terms exactly which behavioral signals and data points it considers to be suspicious.

Jeff Because if it doesn't, it's an undefendable decision, right?

Emily But if an AI is required to explicitly document exactly what it catches and how it catches it, well, how long until highly organised criminal networks start using their own AI models to study those exact explanations, reverse engineer the logic and generate the mathematically perfect undetectable synthetic identity.

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.