EPISODE 25 | Analysis

OSFI 2026 Compliance

JUNE 29, 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 So imagine for just a second that you work in this incredibly high stakes environment where you report directly to two different bosses.

Jeff Oh, man, that's always rough, right?

Emily And the first boss demands that you move at absolute lightning speed. Like if you aren't producing immediate, actionable results, the exact second a problem arises, you are failing.

Jeff Okay. Stressful.

Emily Yeah. But your second boss demands this meticulously documented, painstakingly detailed mathematical explanation for every single move you make.

Jeff Wow.

Emily And if you miss a single step in your documentation, you're also failing. I mean, it sounds like a total nightmare scenario.

Jeff That sounds impossible, honestly.

Emily Exactly. But for Canadian financial institutions in twenty twenty six, this isn't some hypothetical stress dream. It's, uh, it is their everyday operational reality.

Jeff Yeah, it's it's the ultimate operational paradox. You have these huge institutions trying to function at the speed of modern digital finance, while simultaneously being weighed down by the heaviest regulatory anchor we have seen in decades. Right? And the consequences for failing to satisfy either of those demands are just staggering.

Emily Which brings us to our mission for this deep dive. We really want to explore exactly how a Canadian fraud and AML that's anti-money laundering teams are surviving this unprecedented pressure cooker.

Jeff Yeah, it's quite the balancing act.

Emily Yeah, it really is. And we're building this picture today from a couple of highly detailed roadmaps. We have an OSFI 2026 operational update and a comprehensive compliance readiness guide from DataVisor.

Jeff Both really crucial documents right now.

Emily Okay, let's unpack this because even if you don't spend your days worrying about banking compliance, what we're looking at here is a fascinating real time window into how massive complex systems handle the collision of artificial intelligence, intense government oversight, and, well, real time threat detection.

Jeff Yeah. And that collision is completely reshaping the architecture of the entire financial sector. I mean, the technology that moves money has evolved so rapidly that the old ways of governing it, they've simply collapsed.

Emily They just can't keep up.

Jeff Exactly. They couldn't scale.

Emily So let's start with the most visceral change on the ground, the end of what they call the binder era.

Jeff Right? Yeah.

Emily The grace periods for these financial institutions to get their houses in order. They're officially over.

Jeff Long gone.

Emily Right. The office of the Superintendent of Financial Institutions, which we'll just call Osfi for short, has fully activated its integrity and security guideline.

Jeff And that is a huge deal.

Emily It is. It means expanded personnel, background check mandates and this completely new style of examination.

Jeff Yeah. Because for a very long time, the standard procedure for a bank facing an audit was essentially, uh, a theatrical paper exercise.

Emily What do you mean by theatrical?

Jeff Well, you'd have these massive comprehensive policy manuals like literal binders sitting on a dusty shelf somewhere. You'd conduct these scheduled annual reviews. And when the examiners came knocking, you just pointed to the binders and said, look, we are safe because we have a written policy for that.

Emily But the current osfi supervisory cycle has shifted entirely away from that static documentation. Now they want live operational evidence, right? It kind of reminds me of being back in school. You know, you used to have a teacher who would just walk down the aisle, glance at your desk and check to see if you had your homework filled out. Yeah.

Jeff And you could basically fake it.

Emily Exactly. You just had some scribbles on the page. But now that same teacher is standing right over your shoulder, handing you a piece of chalk and making you solve the complex math problem live on the chalkboard.

Jeff And you watch your every move.

Emily Yes, but wait, I have to jump in here because we're talking about massive legacy banking institutions. Some of these places are still running on mainframes built in the nineteen nineties.

Jeff Oh, absolutely.

Emily So how can a bank possibly pivot from doing an annual pay per view to providing live, real time proof of threat detection without entirely halting their day to day operations? Isn't this an impossible engineering standard?

Jeff You're right to be skeptical because they absolutely cannot halt operations. I mean, the financial system never sleeps, right? What's fascinating here is the underlying philosophy driving that change in the server room. Regulators have finally internalized that a static document written six months ago does absolutely nothing to stop a malicious cyber vector that's happening right now.

Emily Yeah, that makes sense.

Jeff So fraud detection infrastructure is no longer just viewed as a neat software tool to save the bank from minor losses. It's being evaluated as a core pillar of operational resilience.

Emily So the examiner doesn't want to read your data classification policy.

Jeff Exactly.

Jeff They want to sit at a terminal and watch how your data streams actively prevent an insider threat in real time as the data is actually flowing.

Emily But again, the legacy systems, right?

Jeff If your security relies on a human analyst remembering to follow a protocol written in a PDF or, you know, manually pulling data from a legacy mainframe to prove a point to an examiner, you've already failed the exam. The only bridge between those legacy systems and this new live evidence standard is intelligent, automated infrastructure. The threats are dynamic, so the proof must be dynamic.

Emily Okay. So if the standard is live operational evidence, we really have to look at who is actually in the room demanding it. Because the bank isn't just dealing with one strict teacher at the chalkboard.

Jeff No they're.

Emily Not. They're dealing with two completely different entities and they have incredibly conflicting demands. Here's where it gets really interesting.

Jeff Oh yeah. This is the crux of it.

Emily The tension between Osfi and Fintrac is the defining structural challenge of twenty twenty six.

Jeff Because both of these watchdogs want a safe, secure financial system. I mean, that's the shared goal, right? But there are specific mandates create a massive operational bottleneck for the banks trying to please both of them simultaneously.

Emily Let's break down the mechanics of this tightrope. On one side, you have osfi. Their mandate is heavily focused on prudential model rigor, risk governance and system explainability.

Jeff Which is a mouthful.

Emily It is. But when we say prudential model rigor, what we basically mean is they want mathematical proof that your system is stable. Yes, they want to know exactly how the engine is built, what variables carry the most weight, and why the engine turns the way it does.

Jeff Right. They want to see the blueprints.

Emily Exactly. But then on the other side you have entry T. They are the tactical side of the house, very tactical. Their focus is tactical dispatch and the immediate effectiveness of suspicious transaction reports or STRs. They want highly contextualized alerts routed to law enforcement immediately.

Jeff Because they need to catch the bad guys before the money actually leaves the country. Right. And if you look closely at how software actually functions, you start to see why optimizing for one of those mandates actively sabotages the other.

Emily How so?

Jeff Well, if a bank tunes its compliance software to satisfy Osfi's demand for rigorous step by step model validation, it introduces a ton of friction into the system.

Emily Oh, because of all the documentation.

Jeff Exactly. It takes processing time and human analyst time to document every single logical pathway in AI took to reach a conclusion, and that friction slows down how fast you can adjudicate an alert.

Emily Okay, but if you decide to prioritize speed, like if you rush those alerts out the door to satisfy fintechs hunger for rapid disclosures, you end up stripping away all the deep analysis.

Jeff Precisely. You submit these thin, context poor reports that basically just say, hey, this transaction looks weird.

Emily And those fail fintechs quality benchmarks anyway because they aren't actionable for law enforcement, right? To use an analogy, it's like being a chef where the first food critic, Osfi, is grading you on the precise, documented molecular science of your recipe.

Jeff I love this analogy.

Emily Yeah, they want a spreadsheet detailing the exact origin and chemical breakdown of every single grain of salt you used. Meanwhile, the second critic, Fintrac, is standing at the exact same table with a stopwatch, demanding immediate delivery of the meal.

Jeff And the chef is completely trapped. Yeah. The reason this dual exposure exists today is because an undefendable AI model is now legally classified as a regulatory liability.

Emily Undefendable, meaning you can't explain it.

Jeff Right? If you deploy a machine learning system to catch fraud, but you cannot instantly demonstrate to an examiner exactly why a specific transaction was flagged or why a different one was cleared to proceed. You are exposed.

Emily So you can't sacrifice explainability for speed, and you can't sacrifice speed for explainability.

Jeff Exactly. The infrastructure has to deliver both simultaneously.

Emily Okay, but if human analysts physically cannot type fast enough to satisfy Fintrac while documenting enough math to satisfy Osfi. The banks are basically forced to rely on artificial intelligence to bridge the gap.

Jeff They have no other choice.

Emily But that brings us to a major reckoning happening right now under guideline E twenty three, which governs model risk management. This guideline essentially spells the death of the black box AI.

Jeff Yeah, guideline E twenty three is arguably the most disruptive piece of regulation to hit data science teams in banking. Oh for sure. Historically, you could just feed an AI millions of past transactions, and it would learn to spot fraud through these hidden patterns. It was a black box. Data goes in, a risk score comes out and nobody really knows what happened in the middle.

Emily But E twenty three.

Emily Makes that illegal.

Jeff Completely. It mandates live model validation and automated drift detection.

Emily Let's spend a minute on data drift, because I think understanding the mechanics of this is crucial.

Jeff It really is.

Emily So a fraud detection model might be highly effective in quarter one, but it can become completely obsolete by quarter three simply because consumer behavior changes, right? For example, let's say the AI was trained to recognize that a fifty dollars e-transfer at two point a m is totally normal because, you know, millions of college students by late night pizza.

Jeff A very common pattern.

Emily Right? So the model learns fifty dollars at two a m equals pizza, clear the transaction. But suddenly a sophisticated new mule network starts using fifty dollars e-transfers at two a m to test stolen accounts.

Jeff And the model just sees the transaction, thinks pizza, and lets the criminals right through.

Emily Because the real.

Emily World reality has drifted away from the historical training data.

Jeff That is the exact mechanism of failure. And under E twenty three, the regulators expect the bank's system to proactively surface that model deterioration before it leads to a massive surge in false negatives.

Emily So the AI has to monitor itself.

Jeff Yes, it has to alert the humans that its own logic is becoming outdated.

Emily Wow. But I think the wildest part of E twenty three is the mandate for alert level explainability.

Jeff Oh, this is a game changer.

Emily The days of an analyst getting a system flag that just says ninety eight percent risk factor are over. They now need automated natural language narrative summaries.

Jeff Which means the system has to generate a plain English paragraph explaining exactly why it calculated that ninety eight percent score based on the specific variables of that exact transaction.

Emily But hold on, I have to push back on this. Okay, good. If I have a machine learning model that successfully blocks ninety eight percent of fraudulent wire transfers, saving the bank millions of dollars, why does anyone care how it reached that conclusion?

Jeff It's a fair question.

Emily Like if it's stopping the bad guys, shouldn't the results matter more than the math behind it?

Jeff Well, it is entirely logical to focus on the success rate, but you have to consider the mechanics of the remaining two percent.

Emily The false positives.

Jeff Exactly. Imagine you are traveling overseas. You try to pay for your hotel, and your legitimate bank account is instantly frozen.

Emily That would be terrifying, right?

Jeff You call the bank panicking, and the analysts on the phone looks at their screen and says, well, the computer gave you a ninety eight percent risk score. You ask what you did wrong, and the analyst replies, I don't know. The computer won't tell us. It just outputs a number. Oh yeah, that's not great.

Emily The sheer panic and lack of recourse in that situation is entirely unacceptable. Transparency is non-negotiable. Automated fraud flags must generate explicit, human readable explanations to satisfy internal teams, to prove the math to external examiners, and ultimately, to protect the consumer from Undefendable system actions.

Jeff That paints a very stark picture. If my account is frozen, the computer said so. Isn't a legal defense.

Emily It definitely isn't. But here is the wrinkle. Banks don't usually build these advanced AI tools from scratch. They rely heavily on specialized fintech partners.

Jeff They outsource a lot of it, right?

Emily They outsource for things like device fingerprinting, analyzing how you hold your phone or type on your keyboard, or biometric screening, and according to guidelines, B10 and B13. Outsourcing the technology does not mean you get to outsource the risk.

Jeff If we connect this to the bigger picture, third party vendor governance has always been a massive vulnerability for large enterprises. Really? Oh, absolutely. A bank might buy a shiny, highly effective AI scoring algorithm from a vendor, plug it into their data stream, and just assume their protected guidelines. B10 and B13 obliterate that assumption.

Emily So the bank cannot claim ignorance about how a third party vendors AI functions.

Jeff Exactly.

Emily Compliance leaders now have to execute what is called an effective challenge. They literally have to force their tech partners to provide transparent model logic using documented, repeatable test scripts.

Jeff It's a huge shift in the power dynamic.

Emily It reminds me of hiring a contractor to build your house. If the roof collapses, the city holds you accountable, not just the contractor. You can't just say, well, the builder said the wood was good.

Jeff It's exactly like plugging a mystery external hard drive into your company's highly secure mainframe.

Emily Oh, yeah.

Jeff You might trust the vendor who handed you the drive, but if a virus on that drive wipes your entire system, the regulator doesn't care who handed it to you. They hold you accountable for plugging it in without knowing exactly what code was on it.

Emily And creating that transparency is incredibly difficult because of the friction it causes with vendors, right?

Jeff Yeah. Because if you are a specialized tech startup, your proprietary algorithm is your entire business model. It's your secret sauce.

Emily Right? So why would a startup let a massive bank tear apart their proprietary algorithm to see how it works just to satisfy osfi?

Jeff They don't want to, but the new guidelines dictate that the bank remains fully liable for model errors. This effective challenge means the bank has to mathematically document exactly how that vendor's algorithm is going to perform under extreme volume spikes, or when faced with evolving unforeseen threat scenarios.

Emily They really have to put it through its paces.

Jeff They have to prove that the external data streams perfectly match their own internal risk appetites. You cannot just accept a vendor's marketing brochure as proof of compliance if the vendor won't open the black box. The bank simply cannot legally use the tool.

Emily Wow. But having the right transparent vendor tools is completely useless if your internal data is fractured.

Jeff This is such an important point.

Emily And that brings us to how these teams actually structure their information to catch bad actors. Because criminals are actively looking for the gaps between systems, they.

Jeff Thrive in the gaps.

Emily Yeah. And the single greatest cause of multi-vector compliance failure today is siloed.

Jeff Data silos are the enemy of intelligence. Historically, fraud teams and AML teams have operated in completely separate universes within the same bank.

Speaker 5 Which is crazy to think about.

Jeff It really is. They use separate databases. They report to different executives. They have completely different mandates.

Emily But the threats don't operate in silos in the sources, they mention things like green fraud and carbon credit manipulation. It sounds like a sci fi thriller. Modern financial crime is so complex, making unified telemetry absolutely vital.

Jeff Let's walk through how a modern criminal actually exploits this. Say a sophisticated mule network executes a rabid account opening online. Okay. A few hours later, they follow it up with a series of structured shell company deposits. A siloed system misses the connection completely.

Speaker 5 Because.

Emily The fraud team only sees the first part.

Jeff Right. The fraud team might just see an account opening that looks slightly anomalous, but it doesn't cross their specific threshold, so they clear it.

Emily And the ML team.

Jeff The AML team, might see a deposit that falls just under the ten thousand dollars reporting threshold. So they clear it. Neither action alone triggers an alarm.

Emily But when you look at them together, the sequence is a massive red flag indicating complex financial crime.

Jeff And that is exactly why unifying those fractured data streams into a single dynamic transaction monitoring layer is vital by combining fraud and AML telemetry into one unified terminal. The institution can detect those complex multi-vector threats before they escalate into systemic failures.

Emily So this is how they solve the dual boss problem.

Jeff Unifying these streams is the exact tactical benefit used to finely balance that osfi and fintrac tightrope we talked about earlier.

Emily Oh I see. If all the data from every department is flowing into one place, the AI can instantly generate the mathematical explanation for osfi and simultaneously package the rich, highly contextualized report for Fintrac at lightning speed.

Jeff Yes, it creates a complete holistic story of the customer's behavior. Instead of just throwing fragmented, meaningless alerts at an exhausted human analyst.

Emily And speaking of the human analyst, Osfi is heavily scrutinizing what they call human in the loop orchestration.

Jeff They're very focused on this.

Emily Because even with all this unified AI, human oversight is still required for complex adjudications, right? But Osfi mandates that any time an analyst or system manager manually overrides a machine generated risk decision, the system must record an indelible, time stamped audit trail.

Jeff Right. Does that mean every single click is cracked?

Emily That's exactly what it means. You can't just click ignore Alert because it's four point five five p m on a Friday and you want to go home. Oh, wow. You have to leave a permanent digital fingerprint explaining exactly why your human intuition disagreed with the machine's mathematical logic.

Jeff That is intense.

Emily It is. Which brings us to the ultimate practical reality check for anyone managing these systems. How do you know if your team has actually achieved this unified, transparent state, or if you were on the brink of a regulatory disaster?

Jeff That's the million dollar question, right?

Emily And the Data Visor guide provides a fifteen point diagnostic test. It's a self-assessment matrix that scores a bank's readiness based on verified automated capabilities versus unvalidated manual workarounds.

Jeff And let me tell you, it is a very sobering metric for a lot of institutions. Yeah. When you tally up the score, it breaks down into three distinct tiers of operational reality.

Emily Okay. What's the first tier?

Jeff So if you score between zero and six you have a critical capability gap. This means your infrastructure relies on unchecked black box models and fractured data silos.

Emily And at that zero to six level, you expose your institution to significant regulatory findings. Massive fines and immediate action is required just to keep the doors open.

Jeff Exactly. The next tier is if you score between seven and twelve. This classifies you as manually, defendable, manually defendable.

Emily So they have the data.

Jeff This means you technically possess the necessary data to satisfy the regulators, but harvesting that data requires massive manual analyst effort.

Emily Oh, I can see where this is going.

Jeff Yeah, the seven to twelve range is where the human cost of poor technology becomes very real. A system in that state relies on extreme stress, massive alert backlogs, and severe staff burnout just to survive a standard multi-week supervisory examination.

Emily I mean, think about your own workplace right now. How often are you relying on manual heroics to get things done? You know, someone staying at their desk until midnight, copying and pasting data from three different legacy spreadsheets just to make a single compliance report work. Because in the world of twenty twenty six, financial compliance, an eight out of fifteen means your team is headed for a wall.

Jeff Oh, they will break under the pressure of a live examination. It is completely unsustainable.

Emily So what's the goal?

Jeff True readiness isn't a state of panic triggered by a regulatory notice. It has to be a standard, effortless daily posture. That is what the top tier represents.

Emily The thirteen to fifteen range, right?

Jeff A score of thirteen to fifteen means you are compliant and automated. You have a low audit risk.

Emily Because at that level, your operational environment naturally surfaces transparent logic data. Yes, you have deployed intelligent infrastructure that is compliant by design. What that means is you have transformed compliance from an agonizing administrative chore into a native, seamless software output.

Jeff The system is doing the heavy lifting of explanation and documentation automatically in the background.

Emily Compliance becomes the natural exhaust fumes of simply doing the job correctly.

Jeff I love that.

Emily Phrasing. So what does this all mean? If we step back and look at the whole picture, Canadian financial compliance in twenty twenty six is an incredibly high stakes balancing act.

Jeff It truly.

Emily Is. On one side, you have osfi demanding, rigorous, mathematically explainable AI testing the absolute molecular science of your server room operations. On the other side, you have an trachea holding the stopwatch, demanding rapid, highly contextualized intelligence to stop criminals in real time.

Jeff And you absolutely cannot rely on static binders sitting on a shelf anymore, right?

Emily And you cannot hide behind proprietary black box algorithms that just output a number and say the computer said so. You have to break down internal data silos, force transparency from your external vendors, and build a system that proves its own logic at the speed of modern finance.

Jeff You know, this raises an important question, one that extends far beyond the server rooms of massive banks. Oh yeah. If these financial institutions are now legally required to deploy AI systems that generate real time, plain language explanations for every single decision they make, proving definitively exactly how and why they reach a conclusion. How long until everyday citizens demand that exact same alert level explainability from the algorithms dictating our healthcare diagnostics, our social media news feeds, or our justice system?

Emily Wow. That is a brilliant thought to leave on. The standard for radical transparency isn't just a banking regulation, it is being forged right now as a blueprint for how humans interact with machines. It's about forcing complex systems to explain themselves in plain language. Thank you for joining us on this deep dive. We hope it gave you some powerful new context. Stay curious. Keep looking behind the curtain and keep questioning the systems working behind the scenes in your own life. 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.