

From Detection to Defensible Outcomes: Rethinking AML Effectiveness for US Banks

Table of Contents

→ Executive Summary	1
→ 1. The Shift in AML: From Detection to Outcomes	2
→ 2. The Detection–Decision Gap	3
→ 3. Why Outcomes Become Inconsistent	4
→ 4. The Growing Importance of SAR Quality	5
→ 5. Regulatory Expectations: Consistency, Explainability, and Auditability	6
→ 6. Moving Toward End-to-End AML Effectiveness	7
→ 7. A Practical Path Forward	8
→ Conclusion	10

“Across the US banking landscape, many institutions continue to rely on legacy, rule-based detection systems that generate high volumes of alerts but fail to consistently identify the most relevant or complex financial crime risk.

As a result, the challenge is not just about processing alerts. It is about ensuring that both detection and decisioning are aligned to risk, and that outcomes are controllable, predictable, and explainable — not dependent on individual interpretation or fragmented processes, under regulatory scrutiny.

Layering automation or investigation on top of weak or fragmented detection does not resolve this challenge, if anything it risks amplifying it.

In today’s environment, effectiveness is defined not by how much you detect, but by how confidently you can resolve and defend risk, ensuring it is fully understood, evidenced, and controlled.”



Brad Levy, CEO, ThetaRay

Executive Summary

Over the past decade, across US banks, detection coverage has expanded — but in many cases, this has been achieved by layering additional rules onto legacy systems rather than fundamentally improving how risk is identified. Most AML programs are not failing at detection, they are failing at decision-making.

As a result, many institutions continue to generate high volumes of alerts without consistently identifying the most relevant or complex financial crime risk.

They continue to face the same challenges: increasing backlogs, inconsistent decision-making, frequent blind spots in identifying meaningful risk, and rising regulatory scrutiny.



“The primary challenge is not the identification of suspicious activity — it is the consistency and defensibility of subsequent decision-making.”

Financial crime risk is not reduced when an alert is generated. It is reduced when that alert is investigated, a decision is made, and that decision can be clearly justified.

Industry data reinforces this reality. **The vast majority of alerts, often as high as 90–95%, are false positives¹**, placing a significant burden on teams. Analysts spend most of their time gathering and reconciling data rather than assessing risk². Regulatory findings continue to highlight delays, inconsistencies, and weaknesses in SAR production as recurring issues³.

Recent enforcement actions have made this shift even clearer. Regulators continue to identify failures in detection, but are increasingly citing weaknesses in inconsistent decision-making, insufficient documentation, and the inability to clearly explain outcomes.

In many cases, the issue is not that risk was missed, but that it could not be defended.

For many institutions, this creates a more fundamental concern: **exposure may already exist within the program but remain invisible until it is tested under regulatory scrutiny.**

The challenge is not only detecting risk – it is ensuring that outcomes are consistent, documented, and defensible at scale.

Ultimately, AML programs are now judged not by how much risk they detect, but by how effectively they manage, evidence, and defend their outcomes

Detection identifies risk.

But regulators evaluate how you resolve it — and how well you can defend it.


The question isn't just whether inconsistency exists but whether it can be identified and addressed before it is exposed through regulatory review.

1. The Shift in AML: From Detection to Outcomes

For many years, AML effectiveness was defined by detection capability. Success was measured by how many alerts were generated, how much activity was monitored, and how quickly alerts could be processed.

That definition is no longer sufficient in today's regulatory environment.

Across the United States, regulatory expectations have evolved. Supervisory bodies such as **FinCEN**, the **OCC**, and the **Federal Reserve** are placing increasing emphasis not on detection alone, but on the quality, consistency, and defensibility of outcomes.



Institutions are now expected to demonstrate not only that suspicious activity was identified, but how it was assessed, how decisions were reached, and whether those decisions can be clearly explained and supported by evidence.

This reflects a fundamental shift in how AML programs are evaluated.

AML effectiveness is defined not by how risk is detected, but by how consistently and defensibly it is resolved.

The question then becomes: if detection capabilities have expanded, why do outcomes remain inconsistent?

2. The Detection–Decision Gap

Detection capabilities have expanded, but often through incremental additions to rule-based systems rather than fundamental improvements in how risk is identified. As a result, many systems continue to generate large volumes of alerts while missing more complex, network-driven, or behavioral patterns of financial crime.

Large financial institutions process millions of alerts each year⁴, the vast majority of which are false positives¹. At the same time, alert complexity has increased, requiring deeper analysis and more contextual understanding.



In practice, this means institutions are investing significant effort to arrive at a relatively small number of defensible outcomes, creating a fundamental imbalance between effort and regulatory value. The result is mounting pressure on teams already operating near capacity.

What emerges is not simply an operational challenge, but a structural imbalance:

Detection is generating more insight than decisioning processes can consistently absorb and defend. In response, many institutions are introducing automation or agentic investigation layers to improve efficiency. However, applying automation to low-quality or inconsistent alerts does not resolve the underlying issue.

If detection is not identifying the most relevant risk, downstream automation can amplify inefficiencies rather than eliminate them.

This imbalance creates bottlenecks and increases the likelihood of inconsistent, non-defensible decisions.

“Across US institutions, a common observation is that detection capabilities have outpaced the evolution of decisioning frameworks, creating challenges in delivering consistent and explainable outcomes at scale.”

Erin Zavalkoff Managing Director, Forensic & Litigation Consulting - Financial Services, FTI Consulting

But value and complexity alone do not explain the problem. The real issue lies in how decisions are made.

3. Why Outcomes Become Inconsistent

The real issue is not simply volume or complexity, it is how decisions are made.

In most institutions, decision-making remains largely unstructured, relying heavily on individual judgment rather than consistent frameworks. Analysts interpret information differently, assemble evidence in different ways, and document conclusions with varying levels of clarity.

This creates variability in outcomes — and that variability becomes regulatory exposure.



“Expanded detection coverage does not, in isolation, reduce risk exposure — it increases reliance on consistent and well-supported downstream decision-making.”

Similar cases can lead to different conclusions. Supporting evidence may be incomplete or difficult to trace. Documentation may lack the structure needed to clearly explain how a decision was reached.

Over time, this becomes measurable. Decision cycles extend. Quality assurance and audit processes require rework. Institutions increase headcount simply to keep pace, rather than addressing the underlying issue.

But the greater risk is less visible.

Inconsistent decisions create exposure that is often only identified during regulatory review, when the cost of remediation is significantly higher. Recent enforcement actions have increasingly cited failures in documentation, decision consistency, and governance, rather than gaps in detection capability, reinforcing that the primary risk now lies in how decisions are made and defended⁵.

By that stage, the issue moves beyond operational — it becomes a matter of regulatory credibility, often requiring extensive remediation and increased supervisory oversight.

Regulatory Insight

In recent supervisory reviews across US banking institutions, it has been identified that similar alerts are not always subject to consistent investigative approaches or decisioning standards.

While suspicious activity had been detected, deficiencies were observed in the documentation and rationale supporting case outcomes. In particular, institutions were unable to consistently demonstrate the basis for escalation or non-escalation decisions across comparable scenarios.

These findings resulted in extended supervisory engagement, requests for additional information, and, in some cases, formal remediation actions focused on strengthening governance, documentation, and decisioning frameworks.

Importantly, the observations did not relate to gaps in detection capability, but to weaknesses in the ability to consistently evidence and support decision outcomes.

In response, some institutions are introducing robotic process automation (RPA) or workflow-driven investigation tools to improve efficiency.

While these approaches can accelerate individual tasks, they do not address the underlying need for consistent, risk-based decisioning.

Without a foundation of high-quality, risk-relevant detection and structured decision frameworks, automation alone cannot ensure defensible outcomes.

4. The Growing Importance of SAR Quality

Suspicious Activity Reports have become one of the most visible indicators of AML program effectiveness.



“The quality of SARs is directly dependent on the consistency and completeness of the underlying decision-making process.”

Regulators are placing increasing emphasis on the clarity, completeness, and consistency of SAR narratives. **FinCEN** guidance explicitly states that SARs must be “clear, complete, and concise⁶”.

Weak or inconsistent decision-making is directly reflected in SARs — through unclear narratives, incomplete evidence, and conclusions that are difficult to follow or defend.

This leads to follow-up inquiries, extended review cycles, and, in some cases, formal remediation requirements.

Improving SAR quality is not simply a reporting challenge.

It is a direct function of how consistently decisions are made and documented upstream.

5. Regulatory Expectations: Consistency, Explainability, and Auditability



“From a supervisory perspective, inconsistency is not solely an efficiency concern — it represents a potential control deficiency.”

Regulatory expectations in the United States are increasingly aligned around three core principles: consistency, explainability, and auditability.

Institutions must ensure that similar cases lead to similar outcomes, regardless of who reviews them. They must clearly explain why decisions were made, what evidence was considered, and how conclusions were reached.

Frameworks such as **SR 11-7** reinforce the importance of transparency, governance, and documentation, particularly in environments where advanced analytics are used⁷.

“Banks must ensure that their BSA/AML compliance programs are reasonably designed to ensure compliance and capable of identifying and reporting suspicious activity.⁸”

FFIEC BSA/AML Examination Manual

For many institutions, the cost of inconsistency is already material, reflected in increased QA overhead, extended review cycles, and regulatory remediation programs.

In many cases, these costs are absorbed as business-as-usual, without addressing the underlying source of inconsistency driving them.

Industry research indicates that investigation and case management processes account for the majority of AML operational cost — meaning inefficiencies in decisioning directly translate into financial and regulatory impact⁹.

6. Moving Toward End-to-End AML Effectiveness

In response, leading US banks are beginning to rethink how AML performance is defined. This shift requires moving beyond static, rules-based detection toward approaches that prioritize risk relevance, behavioral context, and network-level insight.



Risk-based AI models play a critical role in ensuring that alerts are not only generated, but are meaningful, prioritized, and aligned to actual financial crime risk.

Institutions that improve how decisions are made, structured and documented have reported reductions in processing time of up to 40–70 percent¹⁰, alongside meaningful reductions in alert backlogs and improvements in consistency. At the same time, they are increasing investigator capacity without proportional increases in headcount.

The conclusion is clear:

The greatest gains in AML effectiveness now come from improving how decisions are made, evidenced, and defended.


7. A Practical Path Forwards

For most US banks, the question is not whether change is needed — but whether current decisioning processes are sufficient to withstand regulatory scrutiny.

Addressing this challenge does not require wholesale system replacement but it does require a fundamental shift in how decisions are made, structured, and evidenced. This also requires ensuring that improvements in investigation and decisioning are built on a foundation of high-quality, risk-relevant detection, rather than relying on additional layers of automation applied to legacy systems.

Large-scale system replacement programs are complex, disruptive, and subject to heightened scrutiny. As a result, leading institutions are taking a more controlled approach — focusing first on improving how decisions are made and evidenced within their existing environments.

This reflects a broader shift:



“AML transformation is increasingly focused on strengthening decisioning outcomes, rather than solely replacing underlying systems”.

This reflects a broader shift:

1

Standardizing Decision-Making

Institutions are moving away from analyst-dependent processes toward structured, repeatable decision frameworks.

This ensures that:

- **Similar cases lead to similar outcomes**
- **Decisions are supported by consistent criteria**
- **Variability across teams is reduced**

This establishes a consistent control framework across the organization, reducing variability and strengthening defensibility.

2

Structuring Evidence and Documentation

Rather than relying on manually assembled case narratives, leading banks are introducing more structured approaches to how evidence is captured and presented.

This ensures that:

- **Clear traceability from alert to outcome**
- **Stronger audit readiness**
- **Faster, more consistent SAR production**

This ensures that every outcome is traceable, auditable, and aligned with regulatory expectations.

3

Embedding Explainability into the Process

Explainability is no longer an output — it must be built into how decisions are made.

This means ensuring that every outcome can answer:

- **What was identified?**
- **What evidence was considered?**
- **Why was this decision reached?**

This enables institutions to respond confidently to regulatory challenges and supervisory review.

4

Aligning with Regulatory Expectations from the Outset

US regulators are not prescribing specific technologies — but they are clear on expectations.

Guidance from FinCEN, the OCC, and the Federal Reserve consistently emphasizes:

- **Consistency of outcomes**
- **Transparency of decision-making**
- **Full auditability of the process**

Institutions that prioritize these principles are better positioned to demonstrate control — even as their environments evolve.

This shifts AML programs from reactive compliance to demonstrable control.

A Controlled Path to Modernization

By focusing on decisioning first, institutions can improve performance within their current frameworks while establishing a clear path toward longer-term transformation.

This approach allows institutions to reduce immediate exposure while establishing a controlled path toward modernization.



“The most effective transformations are not defined by system replacement, but by the ability to deliver consistent, defensible outcomes at scale.”

The objective is not simply to improve efficiency – it is to ensure that every decision can withstand regulatory scrutiny.

Conclusion:

US banks do not just have a detection problem — they have a consistency and defensibility challenge at scale.

As financial crime becomes more complex and transaction volumes continue to grow, regulatory expectations will continue to increase — particularly around explainability, consistency, and auditability. At the same time, institutions face ongoing pressure to reduce costs while improving effectiveness.

The path forward is not about detecting more risk.

It is about ensuring that every decision is made consistently, supported by clear evidence, and structured in a way that can be easily explained and defended.

Institutions that fail to address this risk will continue to absorb operational cost and regulatory exposure, while those that act will establish a clear advantage in defensibility, efficiency, and regulatory confidence.



Because ultimately, AML effectiveness is not defined by how risk is identified — but by how confidently an institution can stand behind its decisions.

The risk for many institutions is not just inefficiency – it is the possibility that inconsistent or poorly supported decisions may already exist within their programs, only becoming visible under regulatory scrutiny.

Detection identifies potential risk; but supervisory evaluation is based on how that risk is assessed, resolved, and evidenced.

Detection identifies potential risk; but supervisory evaluation is based on how that risk is assessed, resolved, and evidenced.

References

1. McKinsey & Company, A vision for managing financial crime risk; NICE Actimize, AML false positive insights
2. McKinsey & Company, Financial crime: Fighting the war with data; Celent, AML operations research
3. FFIEC, BSA/AML Examination Manual; FinCEN, SAR Activity Review
4. Deloitte, Anti-Money Laundering Survey; KPMG, Financial Crime Outlook
5. Public enforcement actions and consent orders from US regulators including the OCC, Federal Reserve, and FinCEN
6. FinCEN, Suspicious Activity Report (SAR) Filing Instructions
7. Federal Reserve, SR 11-7: Model Risk Management Guidance
8. FFIEC, BSA/AML Examination Manual
9. Celent, AML operations benchmarks; McKinsey & Company AML operating model research
10. Celent and industry benchmarks on AML investigation efficiency improvements

ThetaRay enables banks to connect detection with consistent, explainable decisioning and resolution, helping ensure that every outcome is structured, defensible, and aligned with regulatory expectations.

Visit thetaray.com or email info@thetaray.com

 **THETARAY**