ORACLE

Evaluating and Improving Transaction Monitoring Systems – A Better Way

A better way forward: A holistic and adversarial approach to evaluating transaction monitoring systems improves outcomes and conserves resources

July, 2023, Version [1.0]

Public

ORACLE

Purpose

This document proposes a new approach to evaluate and optimize transaction-monitoring systems and outlines the benefits of this new approach. It also demonstrates how Oracle Financial Services Compliance Agent can help financial institutions adopt this modern approach.

ORACLE

A better way forward: Holistic and adversarial approach to evaluating transaction monitoring system improves outcomes and conserves resources

Transaction monitoring has been a critical component of anti-money laundering (AML) compliance for the last 20 years. To keep pace with rapid changes in financial crime and ever-changing regulations, financial institutions devote tremendous amounts of resources annually to maintaining, optimizing, and enhancing transaction monitoring systems.

The workhorses of transaction monitoring have been rules-based systems, commonly referred to as scenarios. Scenarios are essentially simple if-else statements that trigger an alert if a specified combination of conditions are met, e.g., **if X > a and Y > b or Z > c**. In industry parlance, **X, Y,** and **Z** are parameters, while **a**, **b,** and **c** are thresholds.

Financial institutions can choose to deploy a few scenarios or up to several dozen depending on the size and risk profile of the institution. They may also deploy new scenarios in response to new products that the institution brings to market.

Although, transaction-monitoring systems are widely used and accepted by regulators, the current approach to optimizing and evaluating them has several limitations that can be overcome through the novel use of new technology.

Scenario tuning and its limitations

The foremost challenge in running an effective transaction monitoring system is ensuring that the system is monitoring the right activity. In other words, is the system appropriately tuned? If the thresholds of the scenario are too low, you could end up generating too many false positives. If they are too high, you could end up with false negatives—a risky proposition for most institutions.

Financial institutions have traditionally used Above the Line (ATL) and Below the Line (BTL) testing to evaluate whether a scenario is appropriately tuned. There are several limitations with this approach.

First, each scenario is evaluated and tuned independently. This ignores an important property of the system: scenarios do not operate in isolation. Multiple scenarios interact and overlap to create a monitoring mesh. Just as analyzing individual athletes on a sports team does not reveal the overall quality of the team, evaluating each scenario independently does not necessarily provide a clear picture of the quality of the overall system. For example, 11 conservatively tuned scenarios may not yield a high-quality transaction monitoring systems as there could be blind spots that could be exploited by sophisticated actors.

Instead, how could a bank evaluate and improve a transaction monitoring system holistically?

Second, several institutions use the analysis of ATL data to determine if BTL testing is necessary. If the ATL data indicates an absence of effective alerts near current thresholds, BTL testing is deemed unnecessary.

The absence of effective ATL alerts might be a necessary, but not sufficient, condition for AML risk BTL. When institutions deploy scenarios to provide coverage rather than to detect specific activity, ATL results may not be a good predictor of BTL risk. For example, scenarios to detect terrorist financing or wildlife trafficking may not be always

productive but are necessary to mitigate the risk of such criminal activity occurring. This also does not, however, mean that it's prudent to indiscriminately carry out BTL testing for every scenario given the high cost of doing so.

Is there is more systematic way of assessing BTL risk for a scenario which can then inform the decision to carry out BTL testing?

Third, when conducting BTL testing, institutions determine the threshold to be tested arbitrarily using $5,000 or $10,000 increments. There is simply no reason why increments should be multiples of $5,000 or $10,000 or even $2,450.

Is there a more methodical, explainable way to determine the right BTL threshold that should be tested?

Finally, consider the fact that ATL and BTL tuning is typically carried out using 12 to 18 months of historical data. The issues permeating this historical data are well understood by most institutions. Most of the good alerts or suspicious activity reports (SARs) used to tune a scenario, such as rapid movement of funds, are not a result of activity of interest to the scenario. In fact, these alerts might have been tagged as suspicious due to entirely tangential reasons, such as negative news on the focal entity. Many institutions continue to use this problematic data to tune scenarios because removing them will leave little signal to tune the scenario.

The overwhelming majority of customers at financial institutions are law-abiding citizens. Stringent Know Your Customer (KYC) procedures ensure that individuals or corporations with even a hint of suspicion are denied services or have their banking relationship terminated. This means the historical data used to evaluate transaction monitoring systems are largely from benign customers.

Using historical transaction data to evaluate transaction monitoring systems is akin to evaluating the strength of the financial system in 2008 using data from the boom years preceding the financial crisis. This approach never revealed weaknesses in the system. After the 2008 recession, regulators mandated that institutions carry out more rigorous and frequent stress tests to evaluate the resilience and robustness of the financial system.

When it comes to AML, how can we stress test the system so that its weaknesses become apparent, giving us an opportunity to fix it before a sophisticated money launderer exploits it?

A better way forward

We believe that by evaluating transaction-monitoring systems through a holistic lens and by simulating adversarial actors who might evade the system, we can create a better approach to evaluating and improving transaction-monitoring systems.

ORACLE

Holistic approach

We believe that a better way of evaluating transaction monitoring requires taking a holistic perspective, which considers the entire collection of scenarios and controls that monitor customers instead of evaluating each in isolation. Scenarios interact and overlap in ways that affect the system's performance. Viewing transaction monitoring systems through this holistic lens will:

- Reveal opportunities to retire scenarios and relax thresholds by identifying redundancies
- Alert institutions to gaps in the system that can be fixed by raising thresholds or deploying a new scenario

Consider an institution that has deployed three scenarios—A, B, and C.

Each of these scenarios when evaluated in isolation may be productive. However, Scenario A and Scenario B may be alerting on largely the same set of customers. By evaluating these two scenarios collectively, we can discover an opportunity to raise thresholds for Scenario A so that it only monitors customers that are outside the monitoring range of Scenario B.

Further, consider an institution that wants to monitor a new pattern of suspicious activity. Rather than deploying a single scenario to monitor this activity, the institution can evaluate whether the transaction monitoring system as a collective unit is able to monitor that pattern. If the institution identifies a partial gap, it can make adjustments to the system to plug this gap, such as raising thresholds for Scenario C.

A new scenario—Scenario D—needs to be deployed only when a significant gap cannot be addressed by existing scenarios.

Adversarial approach

An adversarial approach, informed by the concept of "ethical hacking," is an important element of a next-generation approach to evaluating and improving transaction monitoring systems.

Ethical hacking is a process of detecting vulnerabilities in an application, system, or organization's infrastructure that an attacker can use to exploit an individual or organization. They use this process to prevent cyberattacks and security breaches by lawfully hacking into the systems and looking for weak points.

Ethical hackers learn and perform hacking in a professional manner, based on the direction of the client, and later, present a maturity scorecard highlighting their overall risk and vulnerabilities and suggestions to improve. – International Council of E-Commerce Consultants

We believe a transaction monitoring system should be evaluated by determining how effectively it can resist an adversarial money launderer who is seeking to move money through the bank. A robust transaction monitoring system will make it infeasible for the money launderer to move money through the bank in a reasonable length of time without triggering alerts.

ORACLE

A modern approach to evaluating transaction monitoring system effectiveness

A system that takes a holistic, adversarial approach to evaluating the effectiveness of transaction monitoring systems can address the limitations of scenario-tuning approaches used by financial institutions today. Specifically, it should enable the financial institution to simulate a money launderer who can test the transaction monitoring system for gaps, much like how an ethical hacker probes a cybersecurity system for vulnerabilities. Such an agent can probe the entire transaction monitoring system rather than each scenario in isolation. Further, the patterns identified by such an agent can be used to identify real BTL risks for each scenario. They also can inform which scenarios should be subjected to BTL testing and which thresholds BTL should be tested.

Besides addressing the limitations discussed earlier, this new approach highlights the pathways or gaps a potential money launderer could exploit. It also can recommend threshold changes and scenarios that can close these gaps.

If an institution wants to launch a new product, the system will be able to determine the resulting AML risk to the institution. Moreover, it can mitigate this AML risk by recommending threshold changes to existing scenarios or new scenarios to monitor this product.

AI provides answers

Unlike the cybersecurity domain, financial institutions do not have the option to hire ethical money launderers to help evaluate their transaction monitoring system.

How can financial institutions implement an adversarial, intelligent approach to evaluating transaction monitoring systems?

Advances in AI—specifically deep learning and reinforcement learning—have made it possible to solve a problem that has been intractable up to this point. At Oracle, we have built Oracle Financial Services Compliance Agent, a solution that improves scenario evaluation and tuning by using deep reinforcement learning to train an intelligent, adversarial agent.

Oracle Financial Services Compliance Agent assesses the transaction monitoring system of an institution holistically, identifies gaps, recommends changes to address these gaps, and provides an estimate of the operational impact of these changes.

If you are interested in learning more, we are happy to set up a demo to show how Oracle Financial Services Compliance Agent can transform the way you assess and evaluate your transaction monitoring system. Please reach out to financialcrime_ww_grp@oracle.com to schedule a demo.

ORACLE

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle