Why Know Your Transaction (KYT) is a breakthrough innovation for transaction screening.

"Context is worth 80 IQ points."
(**Alan Kay**, Computer Scientist, Turing Award laureate 2003)

When speaking with analysts having to resolve sanction screening alerts, the first complaint you hear is usually about false positives: *"There are too many alerts*", "*The system always returns the same hits*", "*we spend our time accepting hits we already accepted*", "*why can't the system just get smarter?*".

To these analysts, we want to say: "We hear you". And we might have good news for you. Let's dive into a new approach that could really make alert handling better, safer and faster.

First, let's be fair to machines: they actually don't return the *same* alerts (screening systems have become better at eliminating alert duplicates) but they return *similar* alerts, and human brains are widely better than computers at detecting patterns, hence our perception of repetitiveness.

Human brains are widely better than computers at detecting patterns.

So how can we make machines detect patterns in transactions?

The idea of Know Your Transactions (KYT) is to provide operators with additional context about the alert by calculating a score representing to what extent the alerted transaction is "usual" or "unusual". We call this value the K-Score: a K-Score of 0 means that the transaction is very common, while a K-Score of 1 means that the transaction is extremely unusual. The K-Score can in the first instance be used to prioritize alerts (alerts on unusual transactions can be investigated first), but over time, real-time profiling allows the analyst to make more informed and faster decisions.

This is based on the observation that, when looking at the whole set of transactions, there is in fact a substantial lot of similarity between transactions. And indeed, actual data sampling shows many senders have recurring payments (e.g., salaries, mobile subscriptions, utility bills…) and many receivers receive similar types of payments, with amounts and frequencies within a predictable range.

From a data science perspective, this looks like a problem that can be solved using profiling, i.e. capturing past events to identify similar patterns and predict future behavior. Profiling is a very compute-intensive and storage-hungry task, so it is usually done on a subset of data, and in batch, typically yielding results the next day. However, using the cloud's elastic computing power and nearly infinite storage, it can now be performed in real-time on the complete set of data.

How it works in practice

In the last twenty years, transaction screening technology has not evolved much and kept relying on the same approach: the program looks at elements in a transaction possibly matching an entry in a sanctions watchlist and, if a match is found, blocks this transaction pending the operator's review.

KYT uses data from transactions to add context in real-time.

This process is ill-suited for today's transactions world: it does not scale well, cannot easily handle new payment instruments and fails to leverage the broader intelligence hidden in the transactions being screened.

KYT is using data from transactions in real-time to profile all parties involved, and predicts whether a transaction is "usual" or not, adding context to the alert, and allowing the analyst to make more informed and faster decisions. KYT leverages data and looks comprehensively at all transactions, thereby generating much richer insights.

In practice, KYT essentially involves four features:

- *Entity resolution*. Profiling transactions as they flow through the platform is not a trivial task, mainly because transaction formats are still not well standardized and structured. The first issue to solve is party resolution: if John Smith appears in 3 different transactions as "`J.SMITH`", "`:50: /123456 JOHN SMITH`" and "`<UltmtDbtr>SMITH,J</UltmtDbtr>`", it should be recognized as the same person in all cases. The same applies for corporates where different spellings and legal forms should not prevent identifying the party.
- *Anonymization*. There is a strong requirement in the industry for not storing private information, pushed by regulations such as GDPR. Therefore it is needed to create real-time profiles for all transaction parties without compromising on privacy. Using a one-way anonymization function on both the sender and the beneficiary of a transaction, we can create a profile for each [sender→beneficiary] arc as well as sender and beneficiary as separated elements (for peer grouping)

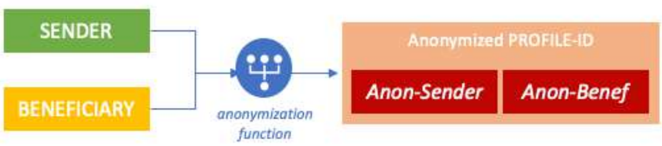

This anonymization technique allows to always update the parties profile, without ever knowing who these parties are.

- *Real-Time Profiling*. Once transaction parties have been identified, the profiling can start by getting as many data points as possible from each transaction. The more dimensions captured, the better the profiling will be at determining whether a match is a potential false positive or not. It is also important to profile every single party participating in a transaction, not just the financial service's customers, as some patterns could only be uncovered by including the full scope. As an example, a notary may not be a client of the service, but many of its clients would transfer money to that notary at a moment in time. If the name of the notary later raises a (false) screening alert, we can start detecting a pattern where each client sending money to the notary raises a similar alert, so this alert is likely a false positive.
- *Scoring*. The profiles are used to create a Machine Learning model trained to detect anomalies and hence predict if the current transaction is "usual" or an outlier. This prediction is providing in real-time as a score that is added to the alert context if the transaction generates a screening hit.

  Over time, the model can learn (unsupervised) what a notary, a car dealer or a utility service is, and what kind of patterns are "expected" for them, even with anonymized profiles where this profession information has been removed.

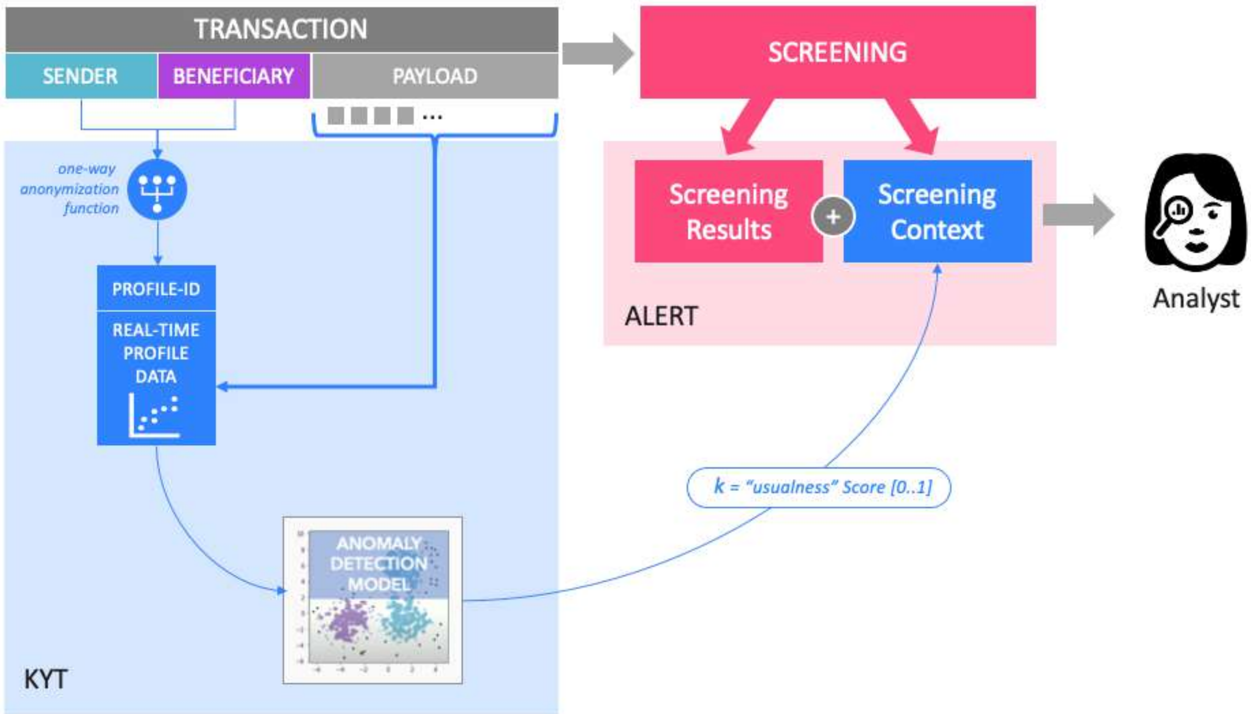

A step-by-step journey

As for any innovation, KYT will need to earn the trust of financial service providers, and therefore adoption will likely go through successive phases:

- In the first phase, KYT will "simply" provide richer context to a screening alert, indicating through a series of analytics scores how much of an outlier a match is. Beyond allowing for better prioritization of alerts, analysts will be able to make sound (and faster) decisions based on the enriched context. Analysts will also contribute to creating a labelled data set containing the alert, prediction and human decision. Such a dataset can in turn be used to train and further refine the model.
- In a second phase, the model would be in a position to suggest a decision that the analyst would simply have to validate or invalidate. Again, this feedback loop would create a dataset allowing the model to be further trained on its decision accuracy.
- Once this cycle is complete, the model would be able to auto-decision most of the alerts with a level of accuracy so high that it could actually be better than its human counterpart. The model would at that point be trusted enough to run on its own, obviously supervised by audit like any other risk function.

In the long term, KYT can drastically reduce manual interventions.

With growing volumes of digital transactions and ever-increasing client expectations for frictionless payments, leveraging old technology paradigms for transaction screening is not an option any longer.

KYT offers a risk-free way to at least fast-track the alert resolution process by providing enriched context to the analysts reviewing the alerts. Longer-term it has the potential to drastically reduce the alerts requiring manual interventions.

Contact Us

info@neterium.io

linkedin.com/company/Neterium

@neterium