Make AML Compliance Easier and Smarter with TigerGraph

Make AML Compliance Easier and Smarter with TigerGraph

EXECUTIVE SUMMARY

Graph is a new way of storing and analyzing data and is especially advantageous for exploring and developing an understanding of any relationships, networks and patterns that may exist in the data set. Therefore, graph naturally lends itself to anti-money laundering ("AML") analysis and investigations, which entails dissecting money flows to identify suspicious linkages.

Within AML compliance, graph can be applied in a wide range of areas, including:

- Enhancing secondary screening of AML alerts to reduce false positive rate.
- Assisting with AML case investigation to improve SAR quality.
- Providing additional insights for customer due diligence ("CDD") and know-your-customer ("KYC") analysis.

In these and many other applications, graph is tremendously powerful in revealing hidden linkages among customers, accounts and transactions through common attributes shared by these entities, such as addresses, phone numbers, social security numbers and IP addresses. These connections would have been much more difficult, if not impossible, for human analysts to identify even with the assistance of conventional analytical software, due to the cognitive limitations of the human brain and computational constraints of traditional relational databases. Once these non-obvious relationships are uncovered by graph, banks can more easily and accurately assess AML risk.

For banks, now is the best time to embark on the journey to add graph into your AML compliance toolkit. The best approach to do so is to take an evolutionary path--starting with employing graph in alert prioritization to reduce false positives, and then gradually deepening and widening the use of graph into case investigation, CDD, and KYC, etc.

TigerGraph differentiates itself from other graph solutions with its ability to scale for massive transaction volume (billions per day) and to perform the deepest analytics. Therefore, TigerGraph can comfortably provide superior graph-assisted transaction monitoring for the largest financial institutions and payment networks on a real-time basis.

AML COMPLIANCE IS HARD BUSINESS

While everyone agrees that AML compliance is very important, few people outside the AML compliance profession fully appreciate how hard it can be to get it right. The main factors contributing to this exceeding difficult task are the following:

- It's expensive. To get it right, AML compliance requires a lot of people, systems, data, and processes to monitor, analyze, and report on suspicious customers, accounts and transactions. These are large investments in the cost center that don't translate into top-line growth.
- It's intellectually challenging. To some extent, finding suspicious transactions is like shooting at a moving target or finding needles in a haystack. Money launderers don't stick to one money-laundering scheme for too long, but instead keep changing their behavior to evade suspicious transaction detection. As a result, AML analysts are often tasked with looking for known unknowns and even unknown unknowns.
- The requirement is stringent. The bar of regulatory requirement is very high and keeps getting higher. On one hand, regulators mandate that banks cannot miss a single productive alert that might warrant a Suspicious Activities Report ("SAR") filing, or in other words, cannot have false negatives. This means banks tune their alert rules to flag anything that is suspicious. Most of these alerts are not tied to real money laundering activities; they are false positives. Meanwhile, banks cannot file SARs on too many alerts either -- cannot engage in defensive SAR filing. This regulatory standard means that banks have to deploy increasingly sophisticated transaction monitoring systems to generate more and more alerts to avoid missing a productive alert, and simultaneously hire increasingly larger teams of analysts to review and investigate all alerts and cases to ensure SAR quality. The analysts filter out as many false positives as possible so that only "truly" suspicious transactions are reported. **The combination of these two efforts leads to a surge in the number of alerts that are generated by transaction monitoring systems but end up not becoming SARs ("false positives"). The industry average false positive rate is over 95%.**
- The consequence of not getting it right is severe. When banks' AML compliance fails to meet regulatory requirements, banks face regulatory enforcement actions including tens-hundreds of millions in fines, restrictions on business activities, reputation damages, and increasingly, compliance officers' and senior executives' personal liabilities.

The work is serious, and the challenge is real. The good news is that graph can help.

HOW GRAPH CAN HELP

Virtually all existing AML compliance systems are built upon relational databases, which store information (customer, account, transaction, etc.) in rows and columns. The relational databases are great tools for indexing and searching for data, as well as for supporting transactions and performing basic statistical analysis; however, the relational databases are poorly-equipped to connect dots and identify relationships, which is essential for analyzing money trails and assessing their AML risk. Using relational databases, in order to find potential connections, analysts need to join a number of tables to run queries. Such queries could take hours or even days to run, rendering any meaningful analysis of linkages among parties and transactions practically impossible.

This is where a graph database can make a big difference. In graph databases, data is stored in nodes – representing customers, entities, and their attributes— and edges – representing connections. When there are multiple nodes and edges, a graph is formed. **Connections between two nodes can be identified by traversals throughout nodes and edges on the graph, which is more efficient and faster than running queries across tables joined together.** Therefore, graph is naturally suitable for showing known relationships and revealing hidden linkages, networks and clusters. **Fund flow trails and parties involved in money laundering activities create a web of entities and relationships through the three phases of money laundering– placement, layering and integration.** Now, this web formed by money flow can be dissected using linkage and network analysis enabled by graph, with the speed and accuracy that was unimaginable before under the relational database. Combined with the conventional AML compliance tools, graph analytics make seemingly fragmented pieces of AML data become a coherent whole and start telling the story behind the money movement.

Below we will walk through a couple of specific use cases to illustrate how graph can improve the effectiveness and efficiency in two areas that contribute most to banks' ever-rising AML compliance cost, namely, transaction monitoring and KYC/CDD. As can be seen from these examples, graph doesn't replace the existing monitoring and analytic tools, but complements them.

Use Case 1 - Enhance Secondary Screening of Alerts to Reduce False Positives

Existing AML transaction monitoring systems produce alerts with a false positive rate above 95%. In order to find those truly risky alerts (less than 5% of the total) on which SARs need to be filed, **banks hire a large number of analysts to conduct reviews and investigations on these alerts. At large banks, this headcount number can run into thousands.** The associated personnel expense is usually one of the biggest contributors to banks' AML compliance cost.

Therefore, the best way for banks to cut cost is to lower the false positive rate without sacrificing the number and quality of SARs filed.

In order to achieve this goal, banks might be tempted to reduce the total number of alerts generated, hoping that this will cut out a lot of false positive alerts. However, this is a dangerous approach, because banks could throw the baby out with the bath water. The regulatory expectation is that banks can't miss a single productive alert, or as the popular saying goes, they need to find the needle in a haystack. So, while reducing the total number of alerts by adjusting the transaction monitoring rule thresholds will certainly reduce the number of false positives, it may cause banks to miss productive alerts as well. This is a sure way to invite regulatory woes. **U.S. Bank was fined $613**

million in February 2018 **by US regulators for lax anti-money laundering control and they aren't alone** - Total US & EU fines since 2009 on banks' misconduct including anti-money laundering violations exceed $342 billion as of September 2017. As a matter of fact, in order to avoid missing that needle in the haystack, banks have been tending to err on the side of caution, driving the total number of alerts generated in the US to steadily increase, rather than decrease, over the past couple of years.

Faced with this flood of alerts, ongoing regulatory scrutiny on SAR quality, and the financial pressure to keep cost in check, banks often turn to prioritization of alerts to improve the efficiency of AML resource allocation, i.e., let more experienced analysts spend more time reviewing higher-risk alerts while leaving lower-risk alerts to less experienced analysts.

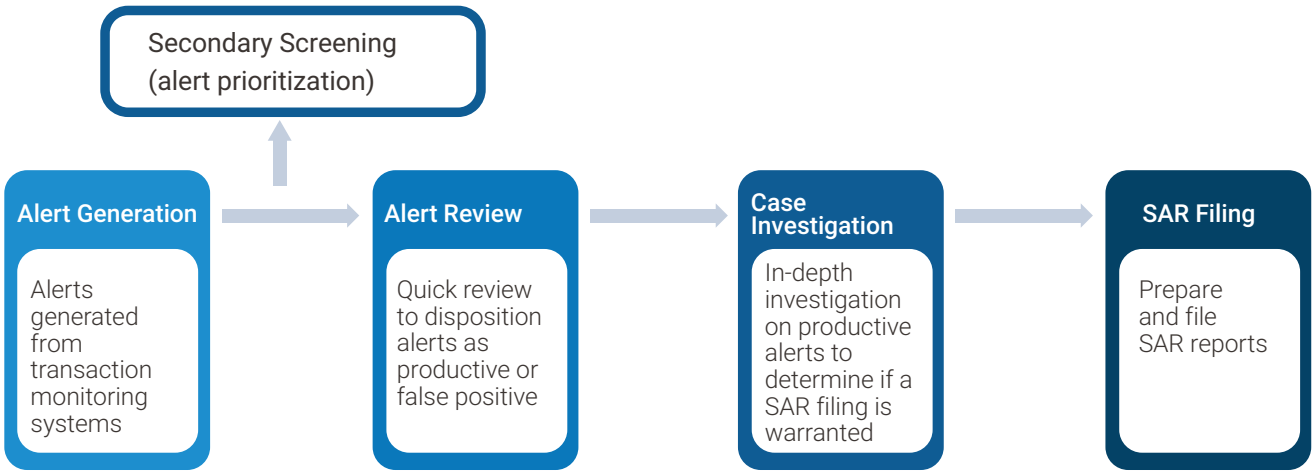

Figure 1 – AML Workflow

As shown in the workflow chart above, after alerts have been generated but prior to the alert review process, many banks conduct a secondary screening to prioritize alerts based on their AML risk. Some banks use a more subjective triage process to differentiate higher-risk alerts from lower-risk ones, while others use quantitative scoring models to achieve this goal. In fact, some leading AML transaction monitoring systems have built-in alert scoring modules that can be customized based on the bank's risk profile. Needless to say, such alert scoring models are more scientific and tend to be more accurate than the subjective approach. Nonetheless, both approaches are invariably built on a relational database and subject to its functional and analytical limitation especially for deep analysis across the network.

Graph can remove this limitation to further improve the efficiency and effectiveness of alert prioritization. **Whereas current alert scoring models generally take the form of multivariate linear formula, graph-based models and analytics add a whole new risk category – network risk, or linkage risk – by revealing hidden links and connections between customers, accounts, transactions and other attributes.**

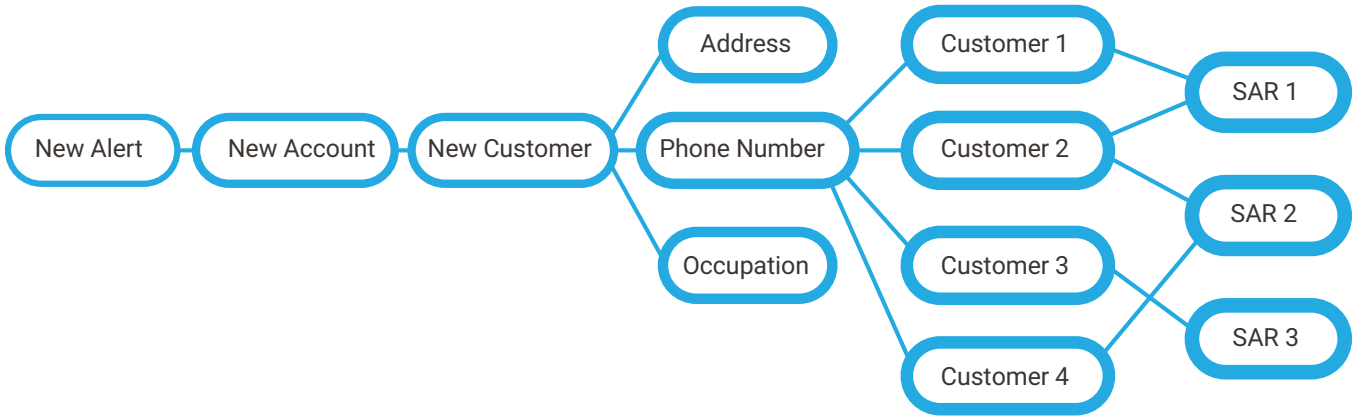

Figure 2 – Uncovering high risk alerts with graph (reducing false negatives)

In the example above, this new alert would likely have been dispositioned as a low-risk alert under a conventional scoring approach, because none of its attributes measured by the traditional scoring model displays any AML risk. However, the conventional approach fails to consider the cluster of high-risk customers that this alert is associated with. This new alert is connected to a new customer, who shares the same phone number with four other high-risk existing customers, on whom multiple SARs have been filed previously. Such a hidden linkage through a phone number would have been quite difficult to uncover using human review or existing models and systems. Thanks to graph, this hidden linkage is revealed, and this new alert would be elevated to a high-risk alert.

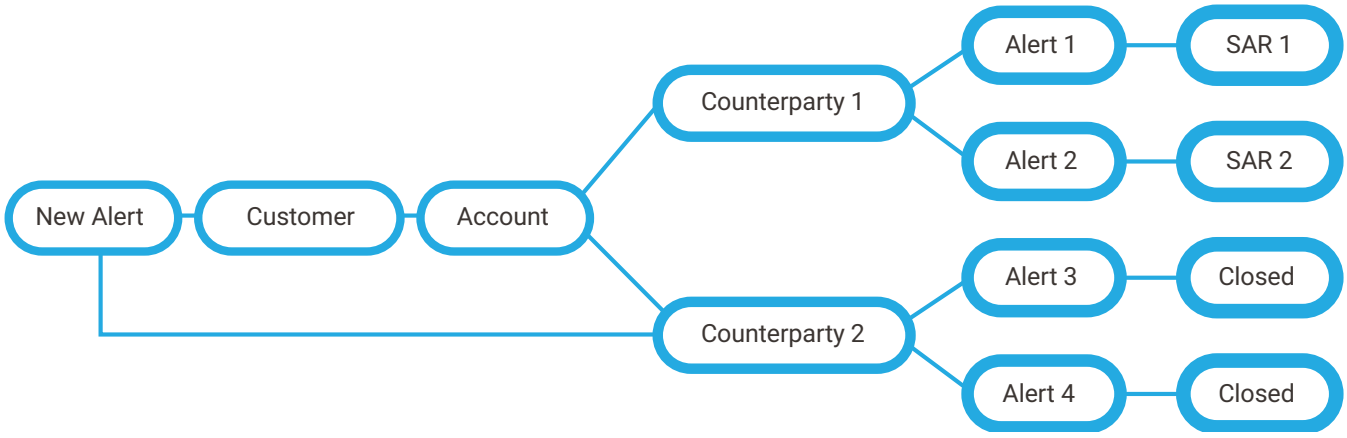

Figure 3 – Ruling out low risk alerts with graph (reducing false positives)

Conversely, in the example above, the traditional approach would probably suggest that this new alert is high risk, because traditional metrics, such as high number of alerts generated and multiple SARs filed on the same account, all point to the direction of high AML risk. Nonetheless, through graph analysis, it turns out that of all the previous alerts, only those related to counterparty 1 became SARs while those related to counterparty 2 have all been closed.

Given that this new alert is related to counterparty 2, it will likely be more similar with alerts 3 and 4, as opposed to alerts 1 and 2, and thus probably should be closed, or at least low risk. **This demonstrates that graph analysis can identify and group "like" alerts where a conventional transaction monitoring system would miss the relationship.**

As can be seen from the two examples above, graph-aided alert prioritization does not replace banks' existing alert review processes, but only enhances it by enabling more complete and accurate AML risk assessment of alerts. This would allow banks to allocate more time and resources to review higher risk alerts, while leaving lower risk alerts to more junior analysts. Some large banks have even started experimenting auto-closure of alerts that fall below certain threshold during this alert-prioritization stage. **This way, banks can save time and lower costs by getting through lower risk alerts quickly, while ensuring the quality of review by having sufficient justification to be confident that no high-risk alerts are missed but all of them receive more careful review in due course.**

Use Case 2 - Assist with Case Investigation to Improve SAR Quality

Graph can also speed up AML case investigators' work while improving its quality. Again, this is made possible by revealing the hidden linkages and transaction patterns through graph analytics aided by visual presentation.

Currently, during case investigation, the analyst would take an in-depth look at both the subject party's as well as key counterparties' historical transactions and due diligence information. This usually translates into pages of account statements or spreadsheets displaying hundreds of (or tens of hundreds of) transactions in chronological order, document after document of KYC and CDD information and banks' internal records. The analyst often needs to log in to multiple systems to retrieve this information and will take notes of her impression and analysis on the case as she reviews these files to form an opinion on the AML risk of the case with little assistance from analytical tools.

This approach is not only inefficient, but also ineffective. Even though each case investigation takes at least two hours to complete, it's still quite difficult for most analysts to understand the full extent of relationships among all the parties and transactions from a naked-eye review of all those documents as separate files, sometimes displayed across different screens.

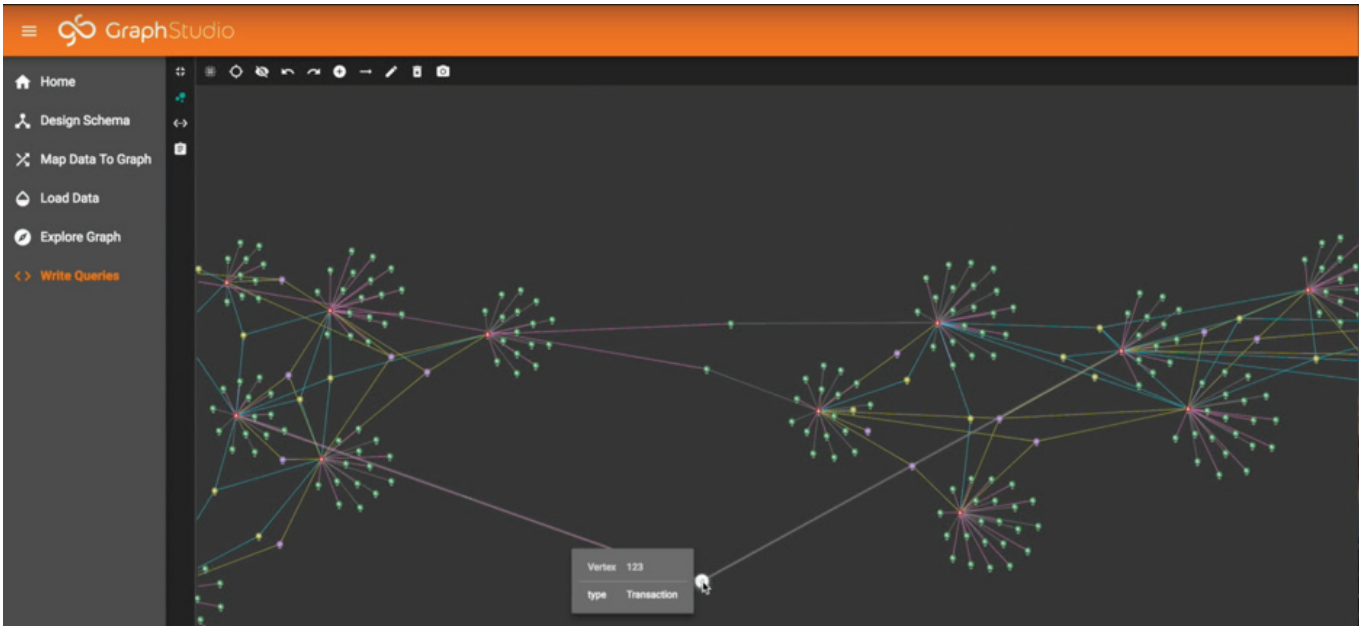

Figure 4 – Visualizing financial network with graph

This is where graph can be of great help. If all the parties, transactions, and due diligence information can be shown in a graph like the above, then the analyst will have a powerful visual aid to comprehend and analyze the transaction pattern and assess the AML risk.

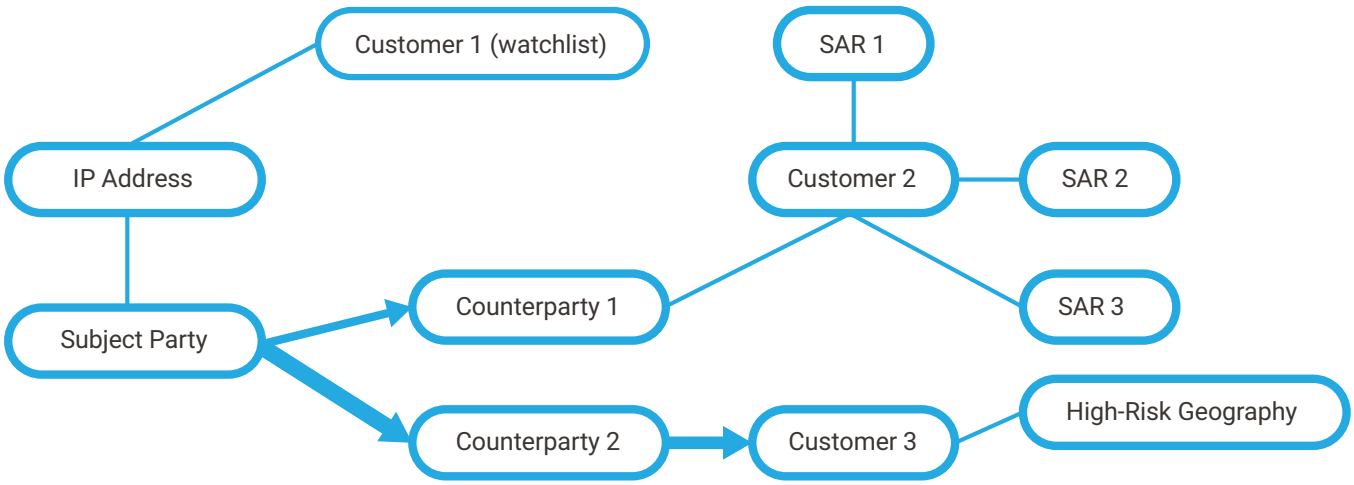

Figure 5 – Visualization and analysis to find hidden relationships with graph

In this example above, thanks to graph analytics and visualization, the analyst can easily see three connections: a) the subject party in the case shares the same IP address with Customer 1 on the bank's watchlist; b) the subject party sent payments to counterparty 1, who in turn is the spouse of Customer 2 on whom multiple SARs have been filed; c) the subject party sent payments to Counterparty 2 who transacts with Customer 3 located in a high-risk geography. All these three linkages would have been difficult for the analyst to uncover by looking at documents and spreadsheets using her naked eyes and processing all the information in her head; however, these hidden relationships become obvious with the assistance of graph.

Again, the graph analytics and visualization tools do not replace the existing analysis tools, but complement them by allowing analysts to look at the same information from a relationship perspective in a more intuitive manner.

Use Case 3 - Upgrade KYC/CDD Research

Similar to the first two use cases related to transaction monitoring, graph can also help banks improve their KYC programs by revealing hidden links and visually presenting the big picture.

Under the current KYC program, banks will conduct due diligence on customers at the onboarding stage as well as on an ongoing basis. Banks often assign a risk rating to each customer to determine the frequency of KYC refresh and adjust these ratings periodically to reflect updated risk profile. Data collected on each customer tends to focus on the static characteristics of customer, such as geography, occupation or industry, types of products used, account age. But this approach will tell banks very little about a customer's risk arising from the linkages he has or the networks he belongs to.

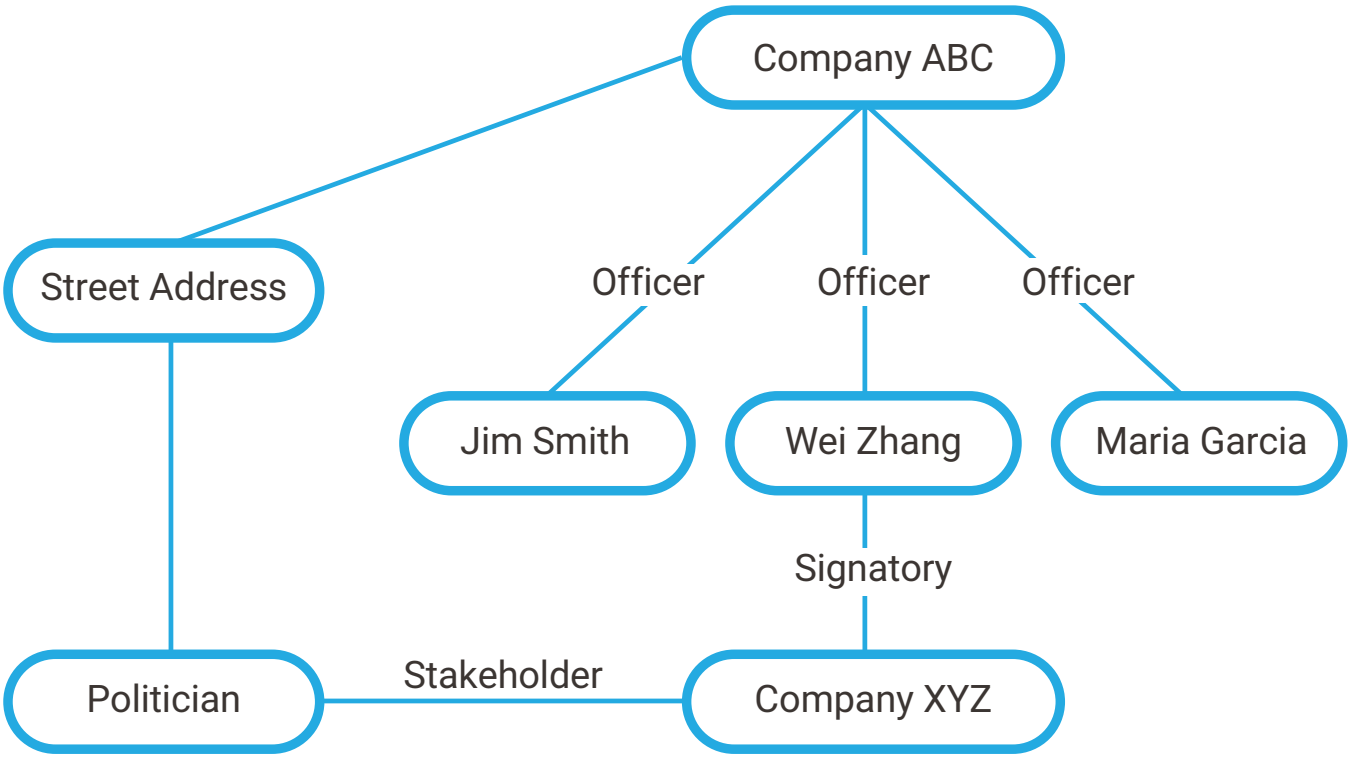

Figure 6 – Applying graph to Customer Due Diligence(CDD) process

Let's walk through an example of an analyst conducting customer due diligence ("CDD") on Company ABC during the customer onboarding process. From the account opening application form, the analyst can see that Company ABC's business address and its three key officers. Standard due diligence performed on the address and officers finds nothing suspicious about the address or the three key officers. Under the traditional approach, the onboarding CDD process would probably end here, and the bank would likely proceed with opening a new account for Company ABC.

However, with the aid of graph analytics, the onboarding CDD on Company ABC can take into consideration more relevant information and reveal previously-unknown risk factors. After Company ABC's information is connected with the bank's internal databases and information collected from third-parties (including commercial database vendors and government websites, etc.), all this information concerning Company ABC forms a large business network graph consisting of nodes and edges. Graph analytics then helps connect Company ABC with the Politician and another Company XYZ of which Politician is a shareholder. Through this linkage analysis, the analyst can make a reasonable inference that the Politician used his address to register Company ABC, and one of his associates managing Company ABC on his behalf is also authorized to manage Company XYZ, of which Politician is also an investor. There is enough suspicion based on which the analyst will make a recommendation that Company ABC should be subject to enhanced due diligence before account opening. This way, the bank might avoid onboarding a customer whose AML risk exceeds the bank's AML risk tolerance. In very much the same manner, graph can help banks identify high-risk customers on an ongoing basis, as banks conduct periodic KYC refresh and as watchlists get updated.

Predictably, the importance of graph analytics will grow significantly as banks and other financial institutions begin to comply with the FinCEN and FINRA CDD rules regarding the identification of beneficial ownership. Graph naturally lends itself to sifting through complex diagrams and analyzing many layers of corporate ownership structure.

For example, below is the corporate ownership structure of Samsung as of June 2014.

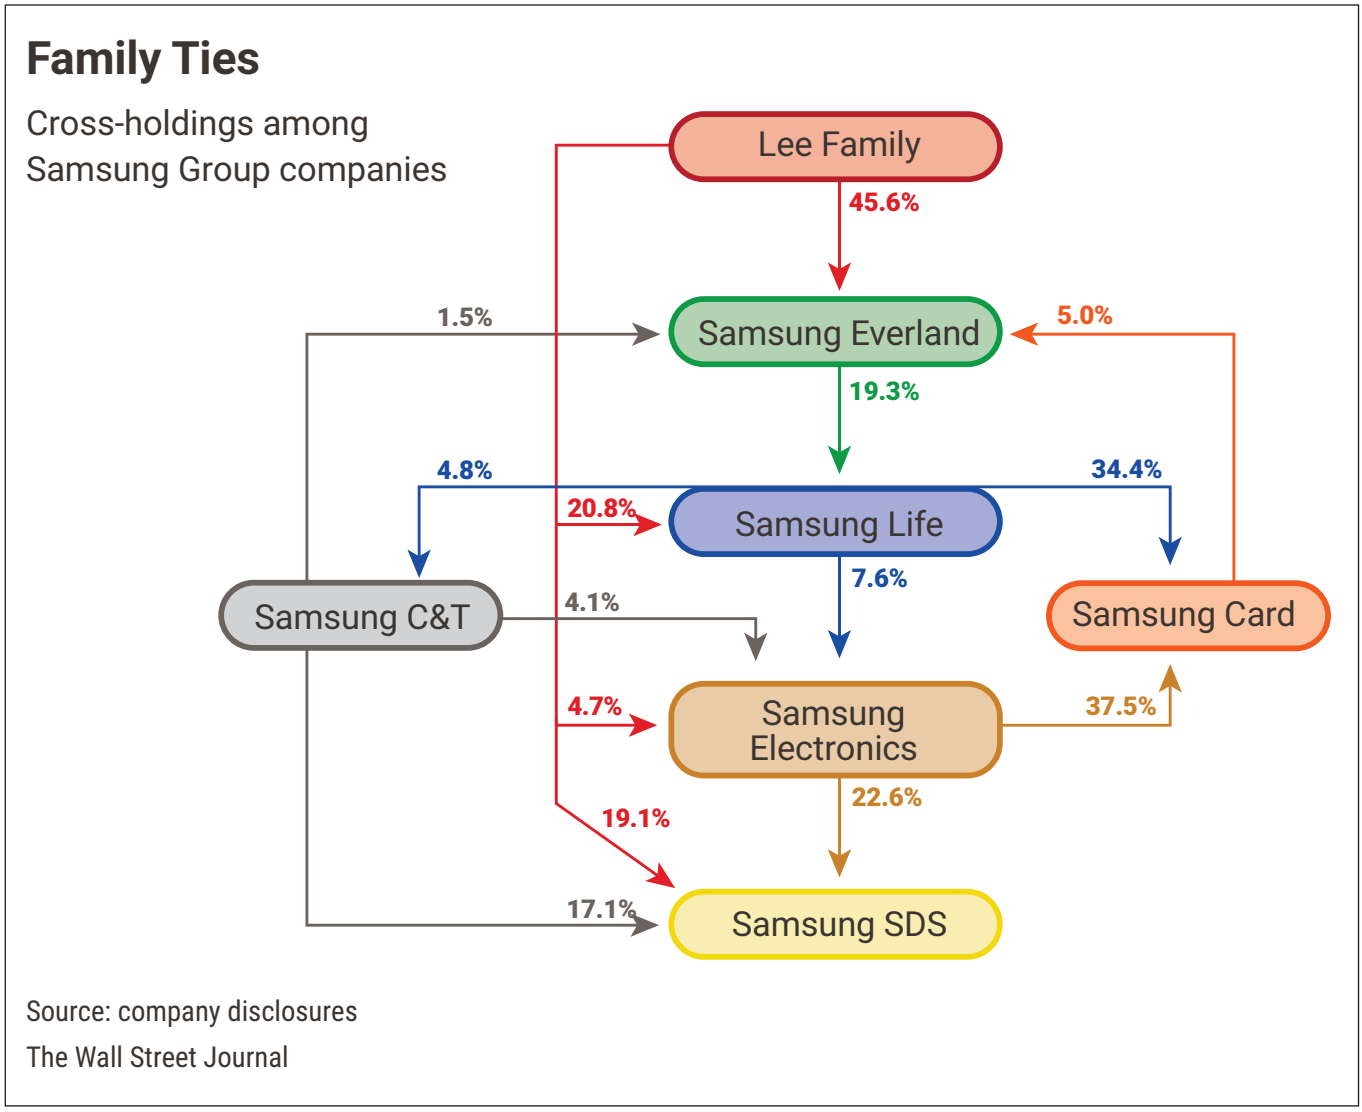

Figure 7 – Samsung corporate structure (June 2014)

If an analyst is trying to determine the beneficial owners of Samsung SDS, it will be much more intuitive and straightforward to map out all the legal entities and complicated ownership structures in a graph similar with the one above. Just imagine how difficult it will be to attempt to analyze this ownership structure by looking at this same information presented in a regular table format.

As banks and other financial institutions begin collecting ultimate beneficial ownership information from new business entity customers and remedying KYC deficiencies from existing customers, graph database can be a very helpful tool for banks to store, update, and analyze corporate shareholding structure to comply with the FinCEN CDD rules.

GETTING STARTED

The three use cases above only offer a glimpse of the opportunities for how graph technology--such as TigerGraph--can significantly improve banks' AML compliance over time. The full potential of such improvement is only bounded by our imagination. As the financial industry is awakening to this huge potential amid continuously heightening regulatory pressure on AML compliance, now is the perfect time to get in on the ground floor of building a new AML compliance toolkit based on real-time graph technology.

Multiple farsighted financial institutions have embarked on this journey and are deploying TigerGraph, a graph database to fulfill certain tasks of their day-to-day AML compliance operations. And the results have been amazing. **For example, Alipay, the online payment subsidiary of Alibaba that processes hundreds of billions of transactions on an annual basis, has been using TigerGraph to monitor and to detect potential fraudulent activities on a real-time basis.** The benefits derived from fraud prevention and operational efficiencies enabled by TigerGraph are enormous. Conceivably, more and more banks of all sizes will follow suit to apply graph analytics to their AML compliance over the coming years.

Embracing a new technology in the highly-regulated banking institutions is always an intricate endeavor, and introducing graph is no exception. Banks need to balance achieving AML compliance enhancement and cost savings through graph technology against IT risks and data governance issues, getting buy-ins from bank's internal stakeholders, managing the regulatory communications and withstanding regulatory scrutiny, and many other challenges.

Recognizing these intertwined complexities, in order to unleash the revolutionary force of graph, financial institutions' best bet is to follow an evolutionary path to let the case for graph become self-evident over time. While a full discussion of the end-to-end project planning and scoping will be saved for another brief, the phased approach depicted below offers a general overview of a thoughtful and methodical approach to adopt graph technology and analytics in the area of AML transaction monitoring in an incremental fashion.

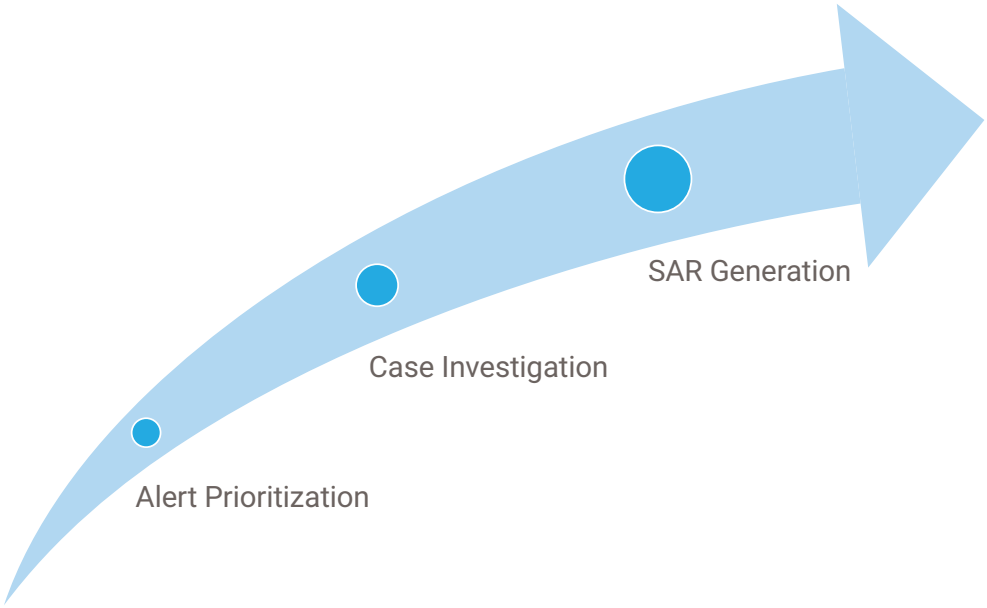

Figure 8 – Transforming AML transaction monitoring process with a three-phase approach

The general philosophy of this approach is that by first taking baby steps to achieve quick wins, banks' internal stakeholders can see tangible benefits in a short period through making relatively small changes to their current processes. As the case for graph builds over time and more internal and external stakeholders get on board, banks can move forward with more expansive use of graph in other AML areas, such as KYC and CDD. This approach ensures that the ecosystem around graph technology and analytics, including IT and data governance, model governance, AML compliance program, senior management oversight, as well as personnel training, compliance testing, audit exam, and regulatory communications, all get the chance to evolve and adapt over time together with graph, so that banks reap the benefit while avoiding pitfalls and unintended consequences.

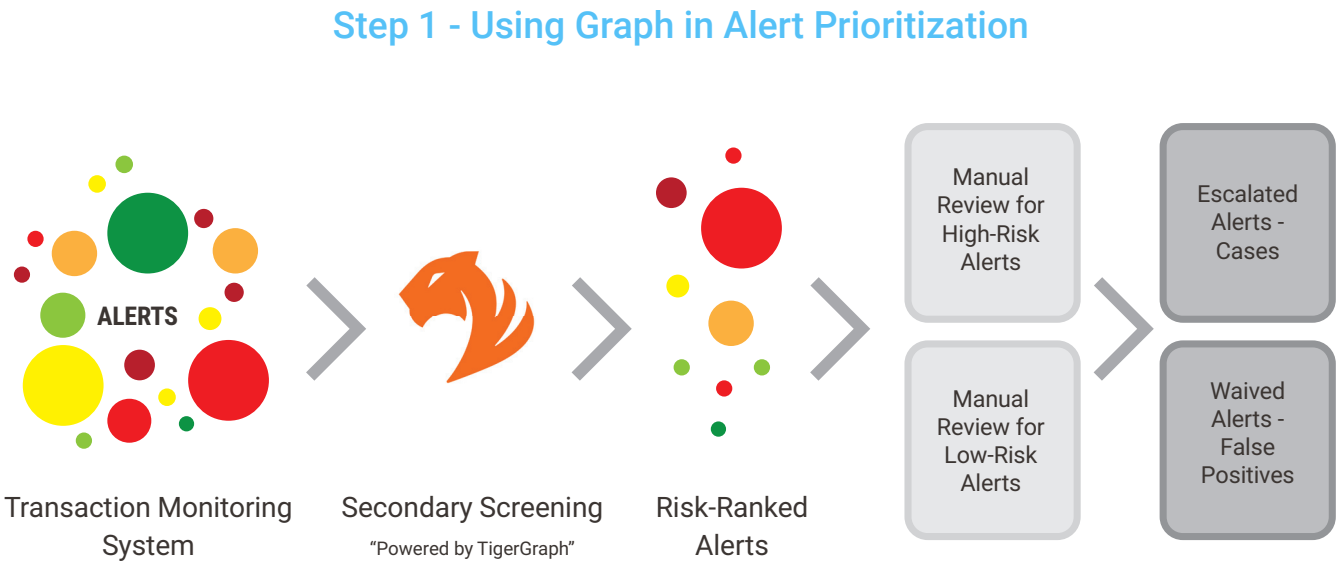

Figure 9 – Reducing false positives with graph for alert prioritization

The best starting point for most banks is to use graph to improve alert prioritization, an area that promises the highest return on investment by trimming the fat from (false positive) alert review.

- This is where the largest money-saving opportunity lies. **When over 95% of the alerts are false positives, even a modest improvement in alert prioritization through better accuracy of AML risk measurement could lead to significantly more efficient allocation of AML resources during the alert triage and review stage, thereby producing sizeable cost-savings.**
- This is an opportunity that banks can act upon immediately. Regardless of which AML transaction monitoring system that a bank currently uses to generate alerts, TigerGraph can always be added to your bank's AML transaction monitoring process as a secondary screening to improve alert prioritization. There will be minimal disruption to existing systems, with no business user interface or new data source required.
- It's quick to deploy and to see measurable results. TigerGraph's plug-and-play approach and agile implementation team ensure that it will take no more than a few weeks to deliver graph-based alert prioritization, and results will be visible almost immediately post implementation.

Step 2 - Applying Graph to Case Investigation

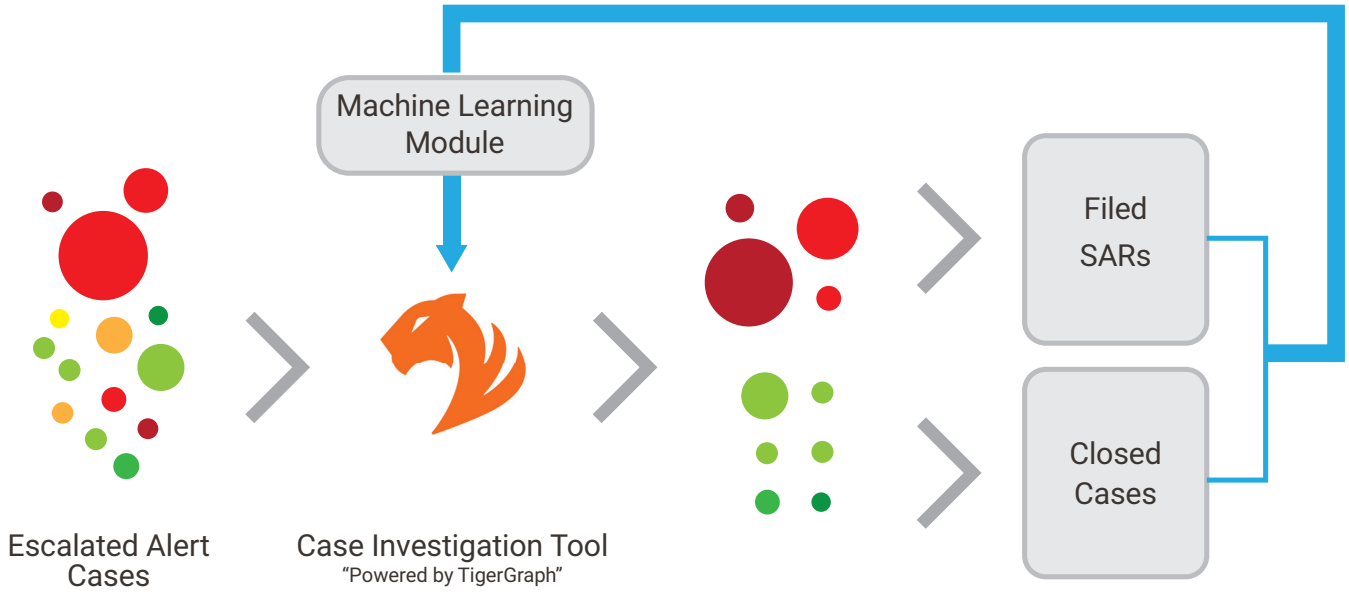

Figure 10 – Improving efficiency and reducing costs for case investigation with graph

After successful implementation of graph for alert prioritization, the logical next step is to apply graph to case investigation.

- The graph tool implemented for alert prioritization will serve as the foundation for case investigation. All the data feeds and alert prioritization algorithms will continue to be essential for case investigation.
- If you apply graph to case investigation, data and IT systems need enhancement to support the graph, for e.g., your graph tool will need data feeds from more internal databases and external commercial databases. Some of these data feeds will need to be real-time. In addition, the graph analytics needs to be supported by an intuitive data visualization tool (such as GraphStudio from TigerGraph) to allow case investigators to make sense of the data.
- Faster and more thorough case investigation will enhance the effectiveness and efficiency of the entire case investigation team, ultimately leading to higher SAR quality and lower personnel and operational expenses.
- Closed cases as well as final SARs can be fed into the machine learning tool along with new features generated by graph analysis to improve the accuracy of future AML alert generation and prioritization.

Step 3 - Embedding Graph within Transaction Monitoring System Detection Scenarios

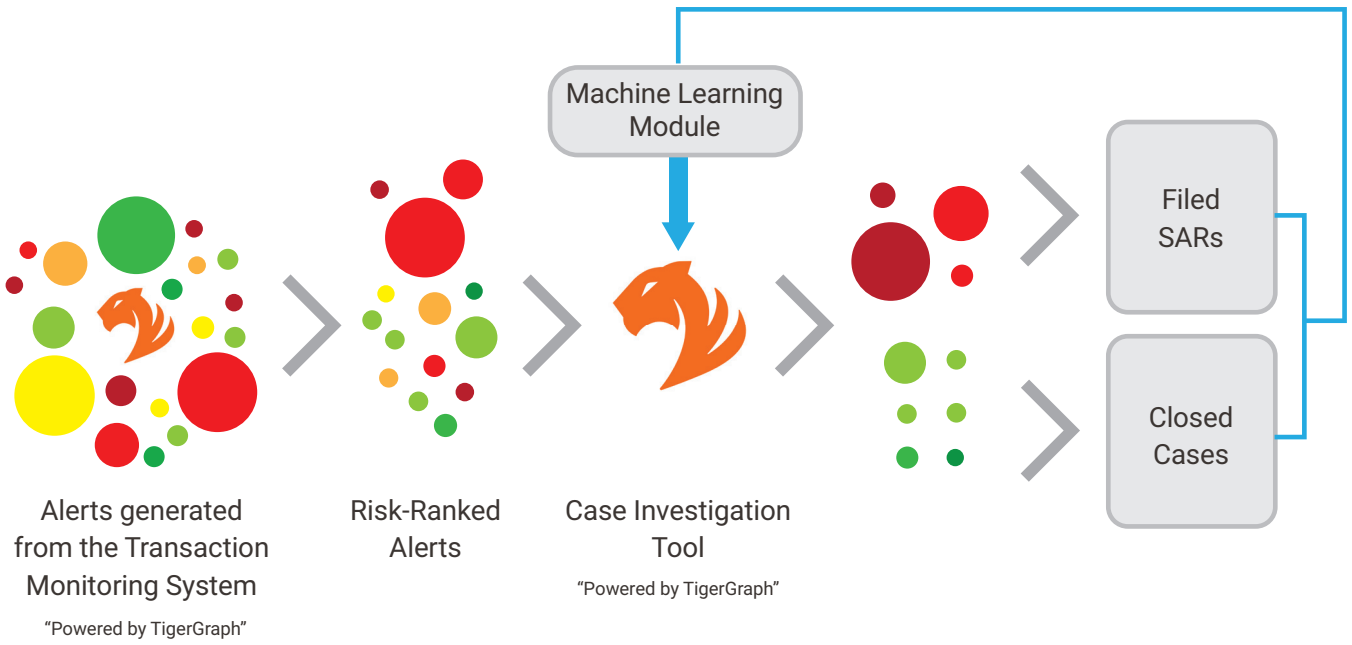

Figure 11 – Transforming end-to-end AML transaction monitoring process with embedded graph solution

The premise of the first two steps is that the existing AML transaction monitoring systems is a given, so the benefits of graph come from efficiency gain and effectiveness enhancement at later stages of alert review and case investigation.

However, if we reconsider that premise, graph can further streamline and improve the AML transaction monitoring when it is incorporated into the AML transaction monitoring system detection scenarios, so that when AML alerts are generated, they will have already reflected AML network risk and linkage risk from the very beginning. There will be no need for second screening based on graph. **Graph-aware transaction monitoring would be tackling the high false positive rate of AML alerts at the root level. We believe that the industry will eventually get there, but we also expect that it will take at least a few years for the ecosystem to evolve before it can support graph-based AML alert generation.**

Across all these three steps and beyond, the benefits of graph can be further magnified if banks employ machine learning algorithms to continuously improve graph analytics. A graph database such as TigerGraph generates new attributes or features based on analysis of connections among payments, accounts and customer data. These new features are highly correlated with detecting AML activity buried deep in intricate networks of accounts and customer entities (false negatives) as well as ruling out false positives. Financial institutions can feed these new graph-based features into the machine learning tool to further improve accuracy of detecting false negatives and reducing false positives.

THE TIGERGRAPH ADVANTAGE

Before we conclude, let's review advantages of TigerGraph's real-time deep link analytics solution for AML compliance. We recognize that there are several other firms providing graph analytics to the financial services industry, and even some of the existing relational-database-powered AML transaction monitoring systems are adding graph analytics modules to their offerings. TigerGraph provides differentiated solution for AML compliance on the following three aspects:

1. TigerGraph is capable of processing hundreds of thousands of payments or financial transactions per second per machine in real-time for AML alert generation and prioritization. This capability makes it possible for TigerGraph to monitor transactions as they come in and issue alerts before transactions get processed and posted. It also enables banks to propagate updated customer information and changes to other business parameters in the system as soon as the updates and changes are received by the bank.
2. TigerGraph digs deeper into the financial network or graph of interconnected financial transactions, accounts and customers. Some risky linkages are deeply hidden. In the world of graph, each direct linkage is called a "hop". So, revealing a deeply-hidden linkage can require many "hops". TigerGraph's deep link analytics can competently identify hidden linkages as many as 10 hops away in real time, thereby exponentially reducing the chance of missing any risky linkages.
3. TigerGraph is designed to handle massive data volume. Currently, TigerGraph is being used by Alipay for 2 billion transactions on a daily basis for Fraud detection and AML compliance. To put things in perspective, according to the Federal Reserve, the total number of non-cash payments in the U.S. in 2015 was 144.1 billion, which amounts to fewer than 0.5 billion transactions per day.

CONCLUSION

To sum it up, real-time graph technology and analytics -- such as TigerGraph--offer a rare opportunity to revolutionize the entire field of AML in banking and financial services. It will become the industry best practice in the near future and may well become part of the regulatory expectations. Implementing any new technology requires some investment and commitment, but at the same time, not having graph in your AML compliance toolkit will likely prove to be even more costly and painful over time. The best way to get started is to take some baby steps (e.g., alert prioritization) to get real benefits first, and then to follow an evolutionary path to more widespread use of graph in AML compliance (e.g., case investigation and KYC/CDD).

As the entire banking and financial services industry is embracing graph in AML compliance, those institutions that start early will get ahead of the curve and turn AML compliance from a burden into their advantage.