WHITE PAPER

How Splunk Can Drive Your Data Strategy in the Fight Against Financial Crime

splunk>
turn data into doing™

Introduction

Fighting financial crime is tough. Building a solution that addresses financial crime can be a lot tougher. The secret to tackling financial crime in the digital age lies in the data, so it's imperative that you develop a data framework that's as comprehensive as cybercriminals are persistent.

Even as you build out a solid framework for operationalizing your data, you're almost sure to encounter significant challenges along the way. For one, financial crime is on the rise and increasing in complexity and style. During difficult economic times, the volume of attacks will increase and techniques will evolve. Professional criminals can be as innovative as the financial firms they target — resorting to data analytics and machine learning to refine their attacks and amplify their volume. While the brute-force attacks of the past still occur, criminals are increasing their sophistication, so firms need to be ready to respond.

Financial crime teams also face a wide range of internal challenges that includes staff in different departments, silos of applications and data, outdated and legacy fraud tools and systems, and a range of structured and unstructured data types that are difficult to compare and analyze. Individual teams can miss out on the opportunity to analyze information in context and often make decisions based on very sparse data. What's more, new products often open the door for new types of financial crime, making it easy for threats to go undetected.

Financial crime teams must adapt to these threats. To combat increasingly sophisticated and evasive financial crime techniques, teams need to collaborate — with each other, and with counterparts in IT and Security — while strategically using relevant data to gain a more complete picture of the financial crime landscape.

That's where Splunk comes in.

Splunk's ability to parse, query and analyze data allows you to tackle all aspects of financial crime, from transactional fraud to money laundering and compliance reporting. Some of the world's largest firms use Splunk as a single platform — helping them break down the silos between the different teams and internal specialists, share data between them, then conduct the necessary compliance reporting and investigations that follow.

In this paper you will learn how to integrate new data sources, unlock value in otherwise discarded data, and bring unprecedented risk modeling capabilities to your enterprise financial crime efforts. Looking closely at specific data types and advanced data analysis techniques, we will outline a six-stage framework that you can implement to bolster your financial crime defenses.

Data challenges create new opportunities for financial institutions

Financial firms face numerous data challenges — both in understanding it and using it effectively — preventing them from realizing its full potential in battling financial crime. However, many of these challenges also present opportunities for firms to use data strategically. Here are a few ways to overcome some of your biggest data hurdles.

Know your data sources

One of the most significant challenges for financial firms is getting a handle on their own data. For one, data comes in various formats, creating complexity. Payments vary in message format by network. Transactions are generally structured, and typically come from a database, mainframe, GL, or from an inbound payment, such as a Swift message. Shared watchlists for money laundering, terror financing, or sanctions usually come from an external source and are often shipped as text files — and usually in different formats.

Data from network traffic, internet banking sites and mobile apps often present the biggest challenge. These logs are verbose, and contain a significant quantity of superfluous text — referred to as digital exhaust — which is often considered to be of little or no perceived value. What's more, they are not well structured and can vary. (For example, when users take different paths through the website.) Thus, attempting to model this data using a database technology designed for structured data is (almost) impossible. Consequently, many firms end up discarding some of the data because they simply don't have the capabilities to analyze it.

However, hidden in this digital exhaust are small nuggets of highly valuable information, which (if they can be found) can lead to valuable insights, and can include:

- Source IP address and location
- Browser version
- Device type
- Failed login attempts
- Pages viewed
- Accounts checked
- Payments made
- Contact center KPIs on account lookup frequency

To find this value, you need to bring your data together, regardless of format. And Splunk can help, with the ability to ingest logs from thousands of systems in real time — correlating relevant data points hidden within those logs. The Data-to-Everything Platform can extract those data points when required, giving them the structure to be analyzed, and making the data fully operational. Your decision makers can detect, resolve and orchestrate successful financial crimes defenses.

The following graphic illustrates how Splunk can help to operationalize data.

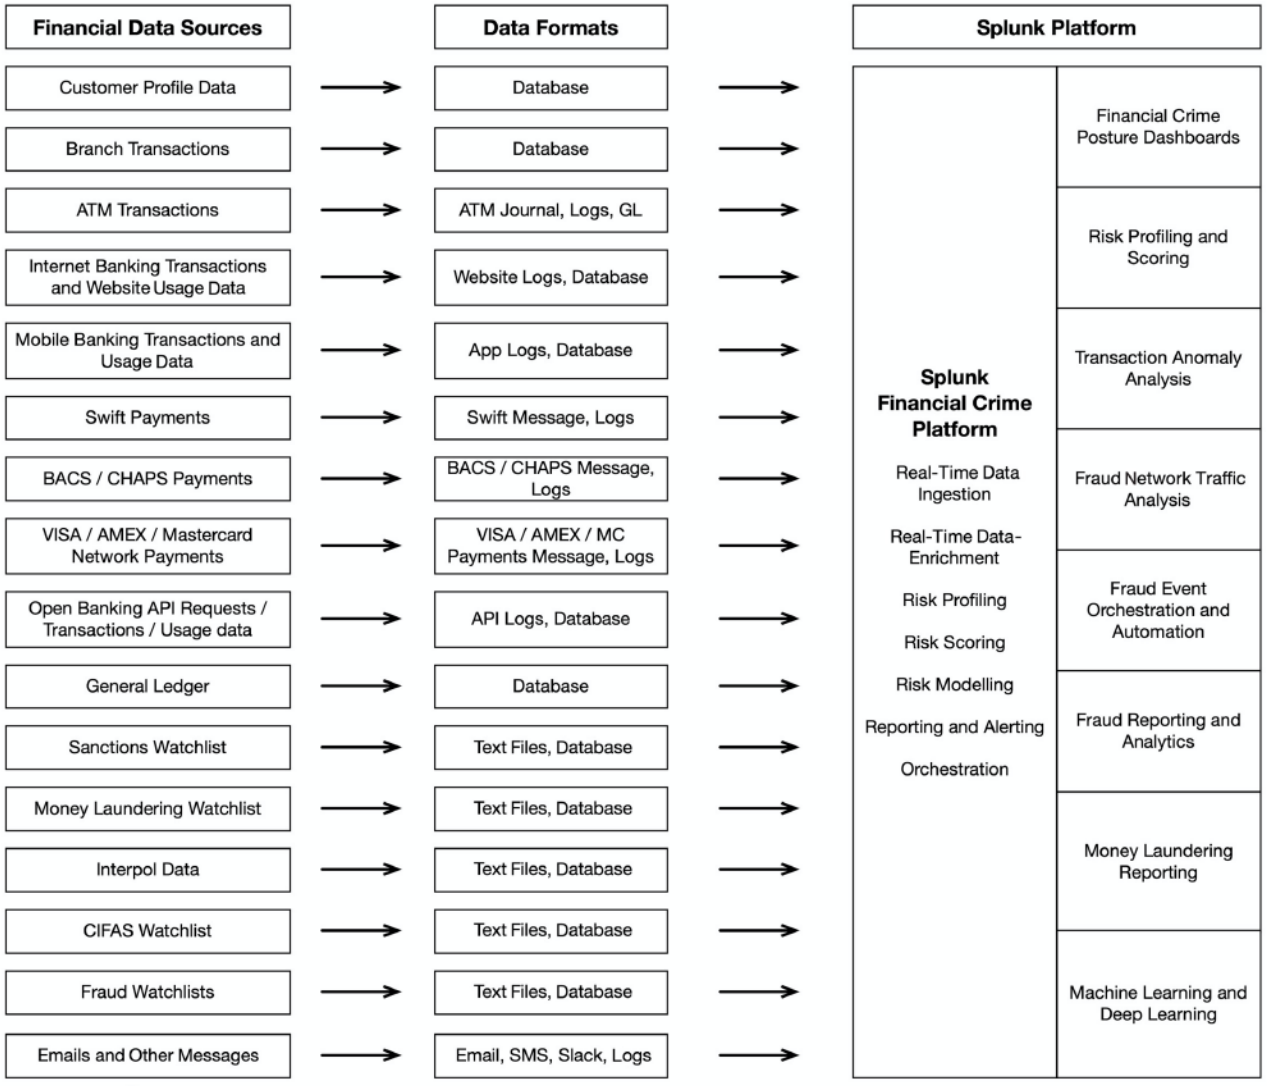

Column 1: A representative list of the many potential sources of data that can be referenced or mined to extract insights needed for effective detection and prevention of financial crime.

Column 2: The likely format for each potential data source.

Column 3: Illustration of how the Splunk Platform can ingest all data types from all data sources to inform the mission-critical functions and capabilities needed to battle financial crime.

Correlate machine data sources

A key roadblock to financial crime fighting often isn't the attackers — rather it's within the structural design of financial crime operations. Historically, firms have organized their operations into silos that address different aspects of financial crime. Most firms already have fraud detection, money laundering detection, and sanctions compliance systems and processes in place, but these areas of defense are often siloed — this is financially inefficient. It also means that the individual teams miss out on the opportunity to analyze information in context with the other processes, while making decisions based on limited data. In short, disconnected teams often create disconnected data, resulting in a host of missed opportunities.

Alternatively, financial firms can realize significant value by pooling their data and resources. As with security systems, outputs from fraud and financial crime systems can be used as inputs to the Splunk platform. By considering transactions alongside internal data on customers and employees, it is possible to look at an enriched view of customers and their transactional activity. Put simply, enriching transactional data leads to higher detection rates, lower false-positives and better outcomes.

Firms can use Splunk to manage sanctions compliance — detecting money laundering activities or collusion, preventing terror-financing and human trafficking, as well as other types of undesirable behavior.

Factor in the human element

Firms can derive significant value from including human behavior data, such as the records of activity on their websites, mobile applications, internal business applications, or systems access. These systems contain tiny clues that can provide meaningful insights to a system or team conducting an investigation. For example, a transaction does not include the number of times a person failed to login, but facts like these are buried in the log files for the website, and can provide the context required to confirm the likelihood of a suspicious transaction.

As a practical example, let's look at a transaction to determine its authenticity. It is possible to bring together the structured world of the transaction — knowing the sender and recipient, the value, the currency and the locations involved. While this is useful, it's not always enough to make a great decision.

Imagine the same transaction, but include additional data points such as the IP address of the sender, the geographical location, the device type, the browser version and the number of times the user tried to log in to their account. We can now determine if the account holder is at home using one of their normal devices, or at a new location, on a new device, using a new browser, and trying out several passwords before successfully logging in — and thus more accurately determine the likelihood of a suspicious transaction.

Taking this one step further, the transaction could reference a list of sanctioned countries or sanctioned companies. It could also reference watchlists of suspected money laundering individuals or states, known criminals or terror financiers. If any of these checks proves positive, the likelihood of a bad transaction increases.

In summary, adding these additional controls leads to a significant improvement on detection and false positive rates, and gives a firm the ability to improve their compliance on money laundering, terror financing and human trafficking regulations. It allows them to set up an operations center where they can monitor all of these threats in combination using a single set of tools and systems, increasing the agility of fraud teams and reducing losses attributed to fraud.

A look at fraud in the contact center

Contact centers are often the weak link in financial crime defenses, often because staff have the ability to look up sensitive data about customers on an ad hoc basis.

Currently contact center employees are given limited authority. Upon employment, people are screened and monitored to make sure that they don't look up too many items of information in a short period of time. The system provides prompts each time they need detailed customer information so that the call center representative can't access the entirety of a customer account all at once.

Yet despite these systems, cybercriminals often execute successful attacks, largely because financial firms are unable to look at the analysis in context with all of the other relevant information on a customer account.

Conversely, the ability to see the analysis of the contact center alongside the transactional history, and the website log analysis would likely provide enough clues to financial firms to prevent these many types of financial crime from occurring. For example, a customer calling in could be using an unusual mobile device, IP address, or calling from a number based in another city. And under normal circumstances, call center representatives would have no way of knowing that the person speaking to them about an account had failed authentication on the website or a mobile device, along with voice or challenge questions.

Splunk, however, can shine a light on criminals' nefarious techniques. Splunk enables contextual analysis that helps identify and flag suspicious outliers, create alerts, and then orchestrate automatic investigations, thus stopping the fraudsters in their tracks.

Implementing a framework to ensure operational data

Fraud specialists typically require both real-time data and a robust library of historical data to support their daily activities and to calibrate models. Splunk offers flexible approaches when it comes to working with data, using data to drive risk indicators, and then using the risk indicators to feed financial crime models, which can weigh individual risk factors, or be aggregated to feed a master model.

The following is a six-stage framework that illustrates Splunk's unique ability to bring in all relevant data sources, automatically organize the data, and correlate across aggregated data, thus overcoming the access and silo issues hampering financial crime efforts today. On this foundation, Splunk enables richer and more meaningful baselining and risk scoring, which in turn enables risk weighting and model development that sets the standard for comprehensive financial crime detection capabilities.

Stage 1: Monitoring events and capturing data

Data can take the form of events on a stream, which can be observed, captured or transformed while still in the stream. These events are then loaded into Splunk in real time, where they are available for search and analysis.

The advantages of being able to observe and modify a stream is that data can be quickly acted upon (before indexing) and modified to include just the critical components in the data pipeline that is loaded into Splunk. This streamlining consumes fewer resources and is faster to process, index and search.

Data can also take the form of logs, metrics, traces, text files, or tabular data from a database, mainframe or other system, including existing fraud point solutions. The ability to cope with all forms of data concurrently, draw insights from each type, and correlate them presents a powerful business advantage.

Stage 2: Data ingestion and field extraction

Splunk can ingest all types of data as described in Stage 1. During the second stage, the data is indexed, enabling Splunk to perform field extraction. Essentially, field extraction enables a piece of unstructured content like a log file to be interrogated as soon as it lands in the index, often just a few seconds after it is created. The process extracts key fields and items of interest that are often buried deep in a log file, and indexes them as a series of events with time as the primary key. This allows meaningful information to be extracted from a log, regardless of what order or form it takes, and is a very powerful capability.

Each journey through a website can be different — determined by the sequence of clicks made by a user — which is rarely predictable. It gives fraud teams critical insights to logins (and failed login attempts), IP addresses, locations, device types, browser versions and other critical business data that can be used to complement a structured transaction. It can also add enough context to the transaction to influence a decision about whether that transaction was authentic or suspicious. Ultimately, that means that teams using Splunk will have a significant advantage when defending against attacks.

Stage 3: Data searching, aggregation and correlation

Once data has been indexed it can be extracted using a search and becomes available for inclusion in a risk indicator.

Searches can be run ad hoc or scheduled. To make them work, Splunk searches organize the data into the appropriate structures in a technique known as 'schema on the fly.' Searches can also aggregate data from multiple places, and perform calculations on the data as it changes.

Correlation searches are frequently used in Splunk solutions, between interesting data-points occurring in multiple, and sometimes, obscure places. Both security and fraud analysts make full use of correlation searches.

The output of these searches, aggregations and correlations can all feed the previously described risk indicators.

Stage 4: Baselining and risk scoring

Risk indicators are used to provide models with an indication of how a specific transaction or entity scores with regard to a specific risk factor. The beauty of this approach is that risk indicators can be fed by diverse data sources, from any system or process. They can be monitored and calibrated independently, and can be based on static or dynamic thresholds.

Risk Indicators need to be baselined, which can be achieved using a variety of techniques. Risk scores can be calculated against static baselines, statistical baselines, machine learning, or even deep learning techniques. Risk indicators can be as simple or as complex as is required to build profiles and searches particular to all financial crime use cases.

Splunk's ability to accommodate diversity in risk indicators is incredibly flexible, enabling a rich library of indicators that can provide warnings in a variety of circumstances. While each risk indicator may only apply to a very discrete or infrequently encountered risk, collectively they contribute to the overall risk weighting.

Stage 5: Risk weighting and modeling

Once the risk indicators have been calculated, they need to be weighted based upon their relevance to the other risk indicators. This is a job for the quants, but the Splunk framework allows it to be done using a simple set of tools and techniques — while also providing continuous monitoring and calibration of the risk weights — so that the model can be tuned for optimal performance. In a production environment, a firm would rely on a model, or set of models designed by its own team of analysts. Risk indicators and models need to be constantly monitored to ensure they are behaving in an expected and appropriate way. The ability to monitor a model allows an analyst to make changes as needed. It also allows them to consistently have a model in production that can be applied to the current market conditions. With Splunk, multiple models can be monitored continuously against a benchmark, ensuring they respond appropriately, and have the ability to be taken out of production when they require calibration.

Stage 6: Model aggregation and calibration

Firms don't just have one model — in fact, it is quite normal for different businesses or regions to have their own unique models. These models can all roll up to a single aggregated model, which can illustrate a company-wide master view of financial crime risk for a large firm.

Built on a wider, comprehensive set of data, these advanced modeling capabilities can be operationalized to drive significant improvement in detection, remediation and reporting. Each unit can create models optimized for their area of focus, from card fraud to money laundering. The thresholds set by these models can trigger automated responses, leveraging Splunk's orchestration capabilities. And all suspicious activity and team responses can be documented end-to-end for reporting purposes.

Summary

Financial firms can experience many gains by bringing financial crime teams together, sharing resources, data and tools, and most importantly working together to design processes and systems for monitoring transactions. Among other things, this allows firms to have common standards from which to work, and define benchmarks for how transactions are scored for risk. It also provides a framework for how regulations such as money laundering, sanctions compliance and insider threats regulations are addressed.

For these teams to effectively collaborate, it is critical to bring together the structured transactional data, customer data and employee data, with the unstructured data from the associated systems such as websites, mobile banking, authentication and other security systems. Unstructured data, which will be new to many teams, can give significant context to transactions that would otherwise go to waste.

Splunk is the engine that brings these disparate worlds together and gives financial firms the upper hand against bad actors. Accessing all data types, organizing them automatically, and correlating them across data sources allows deeper investigation, more refined modeling capability, and end-to-end visibility into events that affect financial crime. Splunk is the platform that powers next-generation financial crime defenses more efficiently and effectively.

To learn more about the financial threat landscape and top challenges to protecting your environment, download a our latest white paper, "On the Hunt for Data: A Look at the Evolving Financial Crime Landscape in the Digital Age."

20-13725-Splunk-Drive Your Data Strategy - Fight Financial Crime-101-WP