Assisting Financial Services Management with Scianta Analytics Extreme Vigilance

Introduction

Complex financial services are difficult enough to monitor and manage without the continual threat of fraud. This hard-to-manage reality is extremely well-suited to Scianta Analytics' Cognitive Computing approach. Seasonality, mobile device access, and continually shifting regulatory environments make it tough to develop sensible monitoring and alerting systems that are flexible enough to handle the realities of FSM. Scianta's Extreme Vigilance products are ideally positioned to assist the financial services support team with solving common business problems.

Elastically scaling service architectures built on virtual machines or containers are changing the landscape of financial services provision, moving the point of monitoring attention away from servers and towards services. Such a transition is even more pronounced for organizations that are building on SaaS, PaaS, or IaaS in the cloud and may not be able to control monitoring directly. This change in monitoring focus necessitates revisiting the prioritization of infrastructure-oriented approaches and increasing the use of behavioral analytics and synthetic monitoring.

In many organizations, the financial services deployment includes an industrial internet of things network of automated teller machines or point of sales systems. These systems communicate events and metrics via standard protocols and data is stored for analysis and alerting purposes, as in standard networks. However, there are sometimes very significant differences to keep in mind. Not only are these systems influenced by swings in customer demand, they can also open to attack by fraudsters. An end-user facing network can be expected to express some of the emergent chaotic behavior of human interaction when outside context problems occur.

Problem Statements

Alerting as a Trailing Indicator

🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨

No one maintaining a Financial Services network wants to see chaos! Perhaps there are known key performance and security indicators to watch, or perhaps they can be discovered. Even so, monitoring dozens to hundreds of indicators for a change that might be important is ridiculously expensive. Avoiding problems costs a lot less than fixing them, which is great if you know about problems before they have occurred. Deviation-from-norm alerts are a good starting point, but it takes a much deeper contextual awareness of transactional patterns and seasonality to avoid predictable alert storms.

Where's the Data?

👓 👓 👓 👓 👓 👀

Data acquisition is also a significant challenge in many FSM environments. While newer equipment might be built to emit signals, a large amount of financial endpoint equipment was never designed for this sort of interaction. Antiquated operating systems, rigid third-party support contracts, and non-negotiable regulatory commitments mean that vulnerabilities and instabilities are built in. Alternate ways can be used to collect data in some organizations, such as passive network capture or access to a database backend managed by another team or vendor; but in other organizations the data scientist can only access extremely limited data sets.

Sound and Fury

🎶🎵 🔥🔥 🔥🔥 🎵🎶

Monitoring is hardly new, and in some cases multiple regimes have come and gone; often they're removed because of the noise generated. While alert noise is arguably better than nothing, a monitoring system that disrupts and exhausts analysts is not helpful. Analysts expect monitoring systems to weigh risk and context appropriately, respond smartly to change, and learn from correction. Above all, the reasons for an alert generation must be clearly discoverable so that logical mistakes can be corrected.

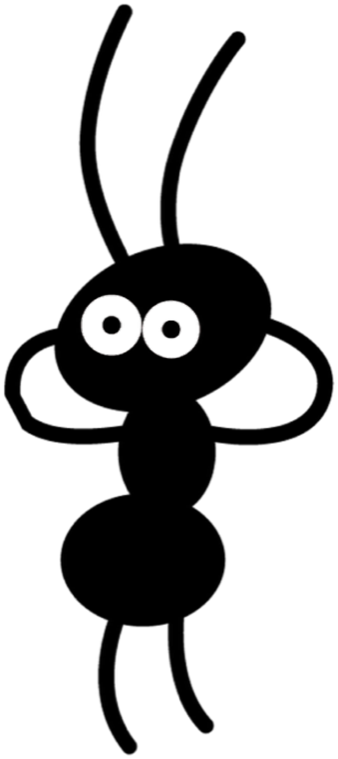

Suggestions

Alerting as a Trailing Indicator

🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨

Extreme Vigilance can model the transactional behavior of an Actor. While metrics and counts may be used to indicate technical problems, alterations in the way that people and systems interact are indications at a business level. Transactional analysis plus anomaly detection opens the door to a better understanding of operations across the board. This approach is particularly valuable for service monitoring, in which a service like credit card validation is the actor and the containers on which it runs are just attributes of an action. By keeping focus on the actor's activities instead of the health of servers, the amount of noise is reduced, and the accuracy of alerting is increased.

Extreme Vigilance also improves on the basic anomaly detection capabilities of more traditional monitoring systems. The analyst and data scientist work together to define a Data Dictionary which describes the events of interest in terms of actors, assets, and actions. These events are reviewed in a Cognitive Model which automatically determines the band of normalcy for each combination of actors and assets by specific time frames. As new values arrive in the data stream, Extreme Vigilance qualitatively measures the fit with observed data, emitting signals when measurements are approaching or breaching calculated thresholds.

These emitted signals can be used to trigger incident alerts of course, but analysts are also able to define crisp rule sets in the Cognitive Rules Engine to describe known compliance issues and take contextual facts into account. For instance, a rule could be written to state that the variable fee charged for a transaction must not drop below 100 dollars. Additionally, rules could be written to express growing concern when the fee is trending downward, or trending downward rapidly, or trending downward very rapidly. Extreme Vigilance calculates reasonable meanings for "very" and "rapidly", while respecting the crisp limit of 100 dollars.

Another excellent way to alert ahead of issues is to use Actor and Peer Analyses. These techniques review recent behavior of resources in order to determine how well the behavior matches with past behavior or the behavior of similar resources. If recent measured values are anomalous, signals are emitted for analysis. Comparing an Actor to itself or to its cohort uncovers subtle variations that may not be visible in threshold anomaly monitoring due to a gradual slope. For instance, peer analysis of PoS registers may be used to uncover a growing product fad, while transactional analysis of wire transfers could be used to find money laundering problems.

Where's the Data?

👓 👓 👓 👓 👓 👀

Missing data is an insurmountable problem for many data analysis systems, leading to accuracy challenging techniques like interpolation and synthetics. Scianta's use of the Splunk platform makes it possible to glean high amounts of value from incomplete, indirect, and disparate data streams. Splunk's rich add-on ecosystem enables access to raw packet capture, infrastructure logs and metrics, database tables, and APIs. While direct access is certainly preferred, Splunk makes it possible to monitor a system indirectly via its impact on infrastructure. In turn, Scianta's behavior analytics can then operate on these indirect signals.

Sound and Fury

🎶🎵 🔥🔥 🔥🔥 🎵🎶

Scianta Extreme Vigilance puts significant effort into producing quality alerts. Signals are generated from a wide variety of rule matches, anomaly detections, and analysis results, but these signals are all weighted. Signal intensity weights are scaled by the severity of the breach, trust level of the model, and criticality of the resources. Analysts can then dial the system's propensity to alert up and down based on their resources and trust in the system's accuracy.

Conclusion

Scianta's Cognitive Computing approach provides transactional analysis and anomaly detection as augmentation to human analysis, purely within the existing Splunk environment. To learn more, see https://www.scianta.com