Detecting Advanced Attackers with Scianta Analytics Extreme Vigilance

Introduction

The modern IT infrastructure is difficult enough to monitor and manage without the continual threat of security breaches. 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 motivated and skilled attackers. Scianta's Extreme Vigilance products are ideally positioned to assist the cybersecurity team with detecting advanced attackers.

Elastically scaling service architectures built on virtual machines or containers are changing the landscape of services provisioning, 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 event collection directly. This change in management focus necessitates revisiting the prioritization of infrastructure-oriented approaches and increasing the use of behavioral analytics and heuristic approaches.

In many organizations, the attack surface includes at least one chaotic, dynamic network. Whether it is semi-controlled corporate endpoints on VPNs, consumers on unmanaged devices, agilely produced and elastically scaling front-ends, or all of the above, security is often an afterthought. With luck these systems may produce events and metrics via standard protocols and with even more luck that data is stored for analysis and alerting purposes. However, access to this data is by no means guaranteed. These systems are an appealing attack surface because they are deeply influenced by outside context problems. Because they express emergent and chaotic behavior, these networks are very difficult to model for abnormal state detection.

Problem Statements

Alerting without Context

🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨

No one responsible for securing an organization 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 and tagging is a significant challenge in many security environments. Between antiquated operating systems, rigid third-party support contracts, and non-negotiable regulatory commitments, security analysts often find 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. Worse, the collected data is difficult to interpret for anyone who is not a highly trained subject matter expert.

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.

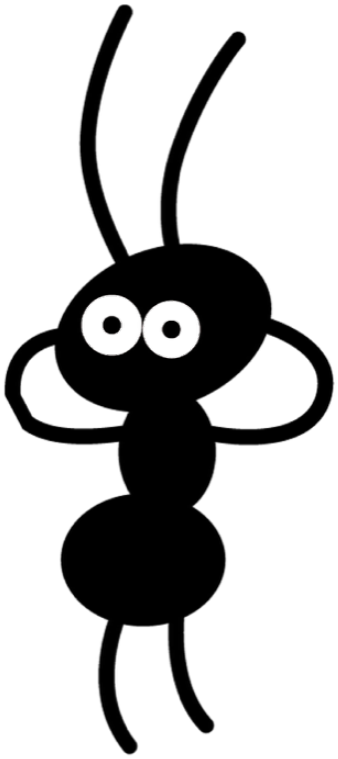

Suggestions

Alerting without Context

🚨 🚨 🚨 🚨 🚨 🚨 🚨 🚨

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 security across the board. This approach is particularly valuable for securing more chaotic edge networks, 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 state of irrelevant attributes, the amount of noise is reduced, and the accuracy of alerting is increased.

Transactional analysis can also be used to detect the context of alerts for risk assessment. Scianta Extreme Vigilance CyberSecurity includes transactional sequences that help analysts recognize when alerts are matching steps in the MITRE ATT&CK framework, so that contextual risk assessment can be done. This approach allows Extreme Vigilance to notify analysts when the threat level associated with a given actor or service indicates advanced attack, freeing analysts from investigating every individual signal in a sea of noise.

While anomaly detection remains problematic for chaotic networks, it is a very useful technique for more predictable systems. Extreme Vigilance improves on the basic anomaly detection capabilities of traditional monitoring. 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 number of recognized intrusion signatures for a network must not rise above 50. Additionally, rules could be written to express growing concern when the count is trending upward, or trending upward rapidly, or trending upward very rapidly. Extreme Vigilance calculates reasonable meanings for “very” and “rapidly”, while respecting the crisp limit of 50 signatures. This Cognitive Rules Engine approach allows alert levels to dynamically scale to the realities of your environment.

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 badge readers may be used to uncover a growing tailgating problem, while transactional analysis of high value auctions could be used to find fraudulent pricing 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. This approach is particularly important for security use cases, in which many different classes of event need to be tagged for many different semantic purposes. For instance, by tagging events by the stage in a kill chain or attack framework, we can determine how well a given set of behaviors fits the model of that framework and determine if an alert is needed.

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