SECURONIX™

Securonix SIEM with Intelligent Analytics

www.securonix.com

Introduction

Security threats today are very different from what they were even five years ago. The security ecosystem itself is multi-dimensional: multi-device (including desktops, tablets, laptops, and mobile devices running Android, iOS, Windows, or Linux), multi-vendor, and multi-host (public cloud, private cloud, on-premises, or hybrid). While this gives enterprises multiple deployment options, it also creates a multitude of new challenges for security analysts.

Every new device that is added to the ecosystem becomes a potential security hole that an attacker could use to gain access to the network. As computing capabilities grow, combined with increasing cloud adoption, attacks continue to become more complex. Today a phishing email can be delivered to your inbox even if you have a secure email gateway, because that gateway relies on source domain filtering to handle malicious email. Attackers today use the same Office 365 or G Suite that your enterprise is using, which means the attacker's email server may be the same one you are using to send an email to your boss.

This is just the tip of the iceberg. Attacks have become complex enough that organizations such as MITRE have created attack frameworks that consolidate information about hundreds of attack patterns in extreme detail. These frameworks are used by companies to actively monitor and secure their environment. Installing a firewall and antivirus is no longer enough protection.

As far back as 2013, the Big Data Working Group of the Cloud Security Alliance published a whitepaper titled "Big Data Analytics for Security Intelligence".[1] In this paper they predicted that the use of big data would rapidly improve security intelligence by reducing the time it takes to correlate and consolidate security data.

However, while they did note that the use of big data in security analytics would improve information security, they did not foresee the changes required in the manner of detection itself. As the way businesses work has changed rapidly over the last few years, so have attacks increased in complexity. The first wave of security information and event management (SIEM) solutions failed to deliver results because they were unable to handle a massive volume of data and because of infrastructure limitations. With today's modern SIEM and security analytics capabilities, however, enterprises can deal much more effectively with high volumes of data as well as complex attacks.

[1]https://cloudsecurityalliance.org/articles/cloud-security-alliance-big-data-working-group-releases-report-on-big-data-analytics-for-security-intelligence/

This paper explores Securonix's security analytics capabilities, focusing on threat chains, behavioral analytics, MITRE-based detection, and the value of having a unified, infinitely scalable solution.

New Security Analytics Capabilities and Threat Landscape

Recent advances in computing technology and models have created two contrasting changes to information security.

1. With the advent of big data and cloud computing, as well as the reduction of storage costs, it is possible to both store extensive volumes of data and sift through and analyze them using big data technology such as Apache Hadoop, NoSQL, Apache Pig, and Apache Hive, or databases such as MongoDB and Apache HBase.
2. At the same time, attacks have become more complex. Attackers are not merely content to try single step attacks, but planned, multi-stage attacks such as in Figure 1.

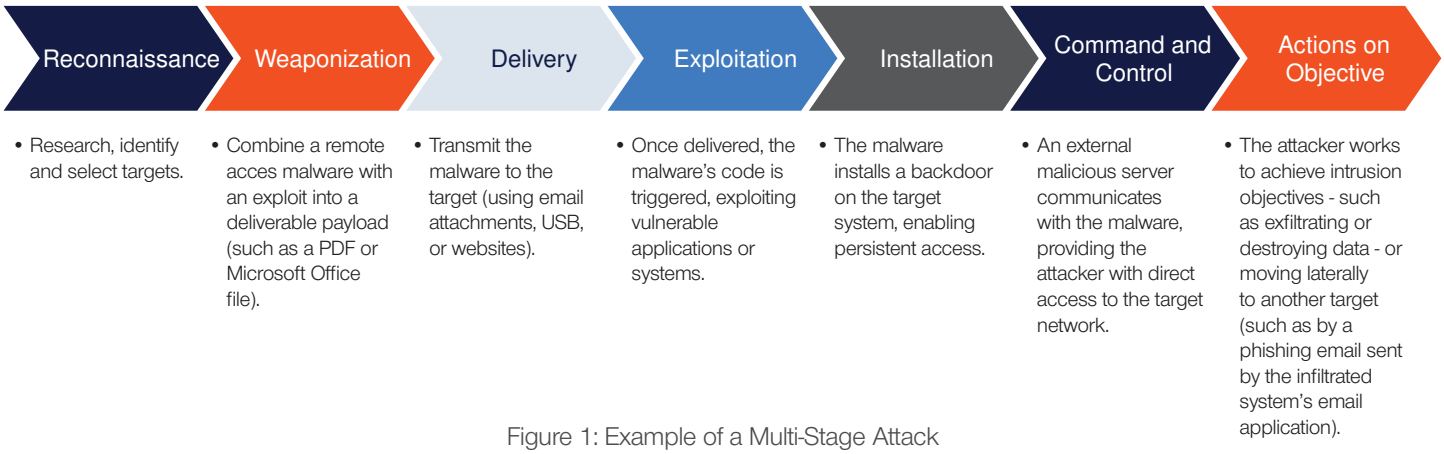

Figure 1: Example of a Multi-Stage Attack

Security teams today deal with more data and more advanced threats, but with the same number of resources. This results in a heavier workload for security professionals, which increases the time required to investigate and respond to advanced threats.

Intelligent Security Analytics

The proliferation of alerts is the most important issue that analysts face today. Aggregating these alerts is now possible with the help of big data tools and techniques that enable the management of huge amounts of data, but analyzing this data is still a huge challenge. Security analysts are, after all, still only human. This is where intelligent, automated security analytics can help reduce the hay that covers the threat needle.

According to the Ponemon Institute in the 2019 Cost of a Data Breach Report from IBM Security, the lifecycle of a malicious attack, from breach to containment, is an average of 314 days. That is more than three-quarters of a year! The report also found that the total cost of a data breach is around $3.92 million USD. However, companies with security automation, such as an artificial intelligence platform, can reduce this cost by around 6%. Using security analytics can also reduce this cost by about a further 5%.[2]

If you chose to add security analytics to your infrastructure, you wouldn't be alone. In Ernst and Young's Global Information Security Survey 2018-2019, cloud computing (52%) and cybersecurity analytics (38%) were the top two priorities for cybersecurity investment, with both seeing an over 50% increase in spending as compared to last year.[3]

Securonix Threat Analytics

Intelligent security analytics can separate the wheat from the chaff for security analysts, using machine learning, heuristic threat detection, and signature-based pattern identification to prioritize actionable alerts.

The Securonix threat detection and analytics framework uses a multi-stage process to reduce noise and provide security operations center (SOC) analysts with a more manageable number of high-risk, actionable alerts that can be acted on immediately. The content of these threat frameworks varies depending on the type of threat, although the pattern of analysis remains the same. These unique threat ecosystems are packaged together as threat chains, which serve as a blueprint to handle threats in that ecosystem. Threat chains exist for multiple use cases, including data exfiltration, account misuse, cyber threat, fraud, medical records, and several others.

First, we will look at the general analysis process, then at an example threat chain.

[2]https://www.ibm.com/security/data-breach
[3]https://www.ey.com/Publication/vwLUAssets/ey-global-information-security-survey-2018-19/$FILE/ey-global-information-security-survey-2018-19.pdf

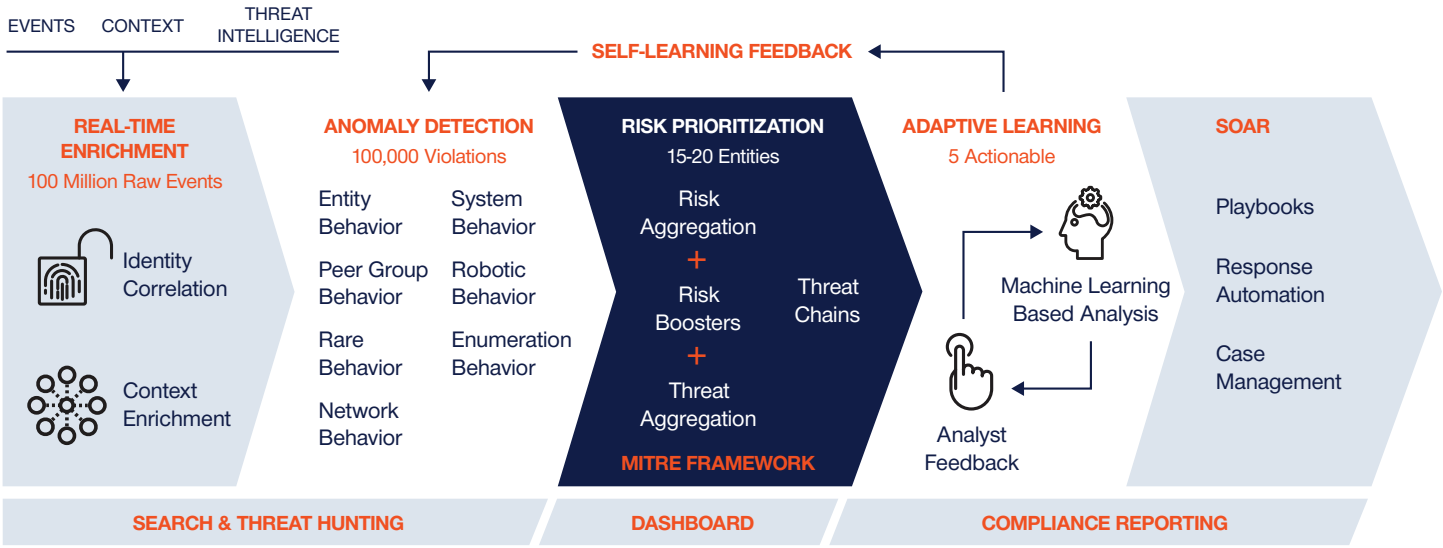

Figure 2: How the Securonix Threat Detection and Analytics Framework Works

The processing of events and alerts takes place through several steps.

Step 1: Real-Time Event Enrichment

Events are combined with relevant contextual information and threat intelligence in real-time – at ingestion – in order to enrich the threat event with additional information that can be used for analytics, risk scoring, and threat investigation. This information allows analysts to clearly articulate and identify the identity, type, and nature of the threat.

In the example shown below, a raw log event has been enriched with contextual and threat intelligence data to convert a difficult to identify threat into a clearly defined one.

The raw event here includes the user ID, action, filename, and IP address. By itself, this log information does not catch the eye. After enrichment, however, the event jumps to a high-risk priority. The entity context information identifies the ID as belonging to an IT admin, who typically would have access rights across the entire IT infrastructure of the company. The IP address' geolocation and threat intelligence data indicate that the access is from a blacklisted IP address in Russia (indicating a case of credential theft) and the accessed asset's contextual information further identifies that sensitive data is being exfiltrated.

Figure 3: A Raw Log Event Is Enriched With Contextual and Threat Intelligence

Step 2: Anomaly Detection With Behavioral Analytics

Post enrichment, the analysis framework looks at multiple behavioral indicators in order to identify anomalous behavior, which may be indicative of a threat developing or already present in the system. These are:

- Entity Behavior: The baseline behavioral pattern of every entity (in this case an entity would refer to a user) is tracked and monitored within the system. This enables the identification of anomalous behavior that typically occurs when there is a threat.
- Peer Group Behavior: The behavior of the entity's peer group (for example system administrators or finance department users, etc.) is baselined and benchmarked against the behavior of individual members. Anomalies, such as a peer group member accessing a file share that is not used by other members of the group, are tagged as possible risks.
- Rare Behavior: Behavior that is rarely seen from the user, such as a significant change in the number of failed logins, could indicate a compromised account.
- Network Behavior: Suspicious network behavior, such as traffic on an unexpected port or related to an unexpected protocol, flowing to a blacklisted or unknown IP address can be indicative of a malware infection.
- System Behavior: Unexpected system behavior, such as open ports that are usually not supposed to be open, beaconing, etc.
- Robotic Behavior: Bot-like behavior, such as multiple login failures within a very short span of time.
- Enumeration Behavior: Behavior patterns that indicate enumeration-based information gathering, such as attempts made to harvest user information from Active Directory (AD) with multiple queries to an AD server.

Behavioral analytics are a critical component of the Securonix analytics framework, using patterns that have been found to correspond to different types of attack attempts. These behavioral analysis detection functions are identified within the Securonix environment as:

- Rare behavior detection
- Frequency/amount spike detection
- Enumeration detection
- Peer outlier detection
- Landspeed detection
- Phishing analyzer
- Network traffic analyzer
- Domain age detection
- Threat intelligence lookup

Behavioral analytics are then combined into threat chains, which link together related behavioral anomalies in order to identify and prioritize low and slow threats that act progressively over a period of time. In figure 4 below, you can see a baseline for the number of failed logins per day for a user. As the chart shows, the average number of failed logins for this user per day is around 5. Yet, on March 1, there were 135 failed logins.

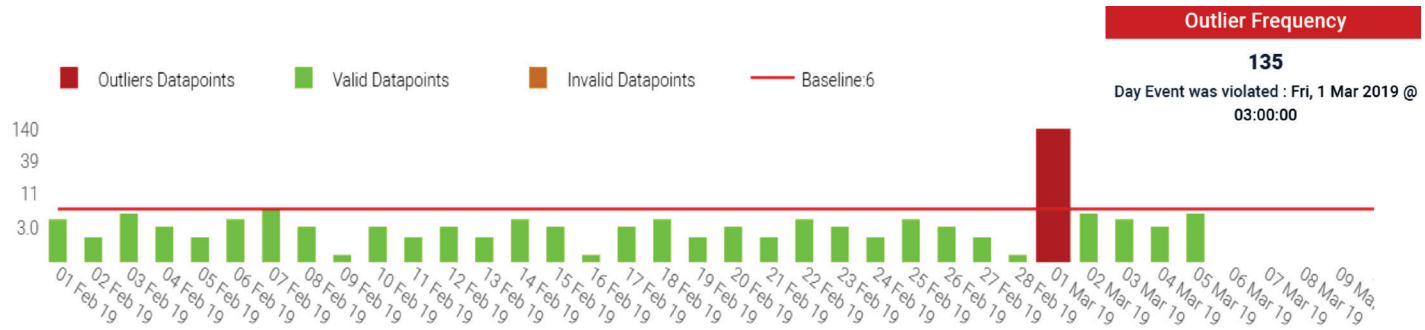

Figure 4: Abnormal Number of Failed Logins Point to Possible Credential Breach

This clearly indicates a deviation from this user's standard behavior, and points to a possible credential misuse. Events such as these are difficult for signature-based solutions to identify, requiring behavioral security analytics to detect.

Behavioral indicators vary for different use cases, however, and the same set of patterns generally cannot be used to detect different types of attacks. Securonix has identified patterns relevant to each type of attack and has used them to create threat chains. One example of behavioral analytics in action is for detecting credit card or ATM card transaction fraud. Securonix provides a complete threat chain package for fraud, as well as other key enterprise use cases.

Step 3: Risk Prioritization and Aggregation With Risk Scoring and Threat Chains

Threat Chains With Behavioral Analytics

Threat chains are a relatively new concept within threat modeling. Securonix threat chains stitch together related alerts, including indicators of compromise (IOC) and tactics, techniques, and procedures (TTP) over a period of time in order to detect patterns of advanced attacks. Individually these indicators may seem like low risk noise, but when correlated to an entity and monitored over time, they can be an indication of an advanced threat. This could be an insider that is collecting data with the intent to exfiltrate it, or an external entity that uses phishing campaigns to hook an unsuspicious user and gain access to the organization.

Securonix threat chains reduce false positives by up to 90% and the cost of operating your SIEM by over 50%.

Threat Chain Content

Securonix threat chains are based on industry standard frameworks such as the MITRE[4] ATT&CK framework and US-CERT[5] Cybersecurity Framework, among others. An example of a threat chain is shown below. It shows a series of indicators starting with a spam email and ending with a system compromise.

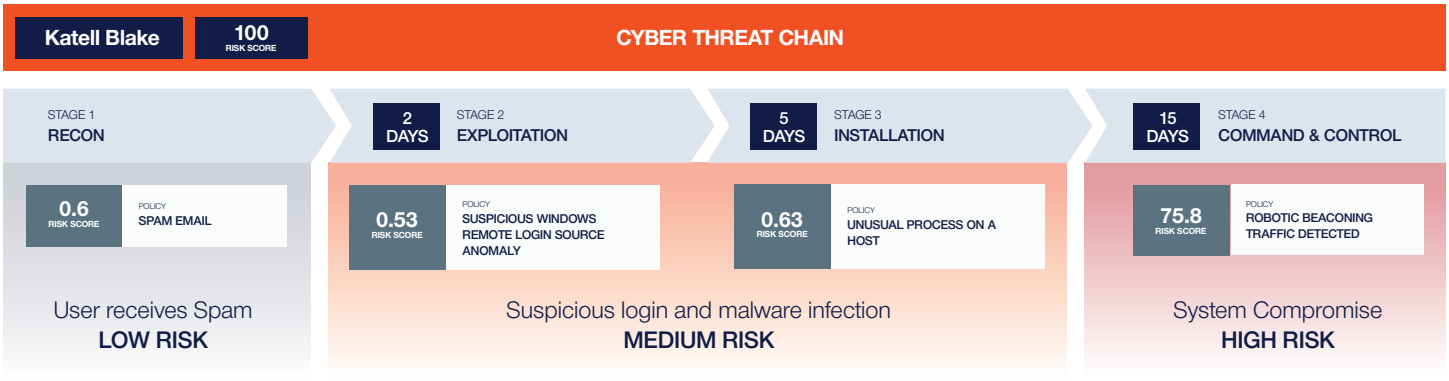

Figure 5: An Example of a Cyber Threat Chain - From a Spam Email to System Compromise

[4]https://www.mitre.org/
[5]https://www.us-cert.gov/resources/cybersecurity-framework

A very important element of threat chains is the ability to elevate the risk score of the entity. There are two risk scoring schemes customers can choose from:

- Static Risk Scoring: A pre-defined static risk score is assigned to the entity upon triggering the threat model. The advantage of this approach is you can set a maximum score relative to your organization to decide escalation and remediation steps.
- Exponential Risk Scoring: In this risk scoring scheme, the risk score increases exponentially as each stage of the threat model is triggered. This ensures that the highest risk threat will always score very high. In a complex IT environment, it might be hard to predict a static high score, in which case exponential scoring is ideal.

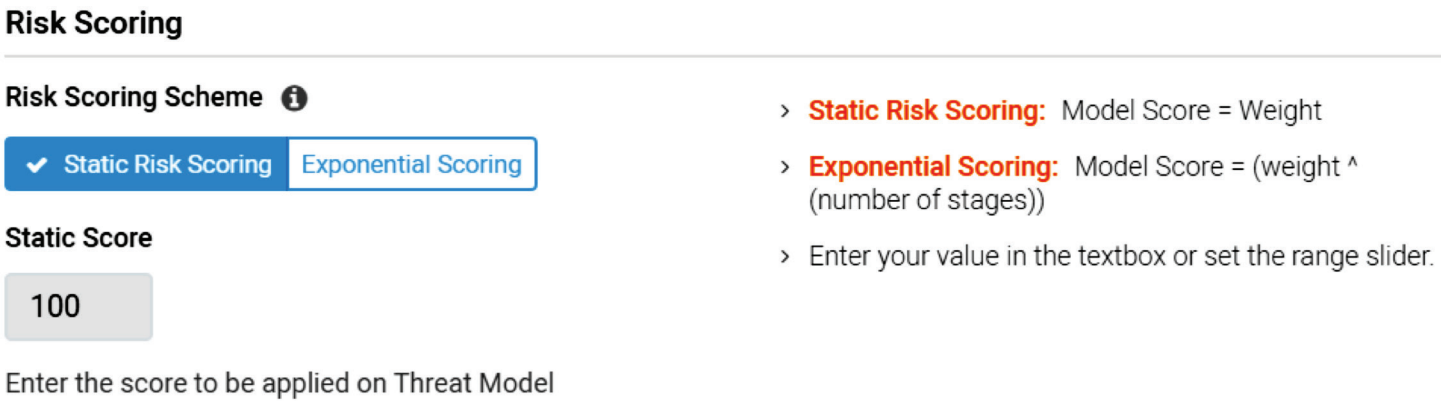

Figure 6: An Example of How Different Risk Scoring Schemes Work

Risk Aggregation

When there are multiple violation alerts linked to the same entity, the risk scoring engine considers all the alerts when scoring the relative risk of the entity. Therefore, if an entity has 5 distinct violations across different entity types (for example, user, IP address, and host) the risk score will aggregate the scores for all five, giving a more holistic threat posture.

Risk Boosters

Data from several technical and non-technical sources provide variables that can be used as risk boosters. Risk boosters modulate the level of risk assigned to an entity based upon parameters such as presence on a watchlist, tagging by a threat intelligence input, or through lookup tables. By adding context from all available resources, risk scoring is both more accurate as well as comprehensive.

Step 4: Adaptive Learning and Automated Self-Tuning

By learning from analyst actions, the Securonix platform enables automated response. Securonix has deployed and matured a self-tuning capability based on adaptive learning for threat snoozing and amplification. The capability uses machine learning techniques such as natural language classification and spatiotemporal analysis to study the patterns of threats that have been detected, the response actions taken by analysts (whether to snooze, suppress, or amplify alerts), as well as other factors such as multiple offenses, offense by multiple actors, rarity context based on the profile of the actor, the victim, etc.

This self-tuning system reduces the number of alerts analysts must deal with and helps customers prioritize threats that need immediate action. It increases efficiency by reducing the overhead on SIEM analysts to configure and tune use cases.

Packaged Content

Securonix packages and delivers analytics content – including connectors, behavior models, threat chains, dashboards, and reports – in the form of apps available in a threat library. This enables organizations to quickly deploy and update content without need of an internal SIEM SME or vendor support.

The Securonix threat library is updated through the Securonix threat exchange portal, which is continuously updated and maintained by the Securonix threat research team. For SaaS customers this content is auto updated on an ongoing basis, ensuring that customers always have the latest and greatest detection in place.

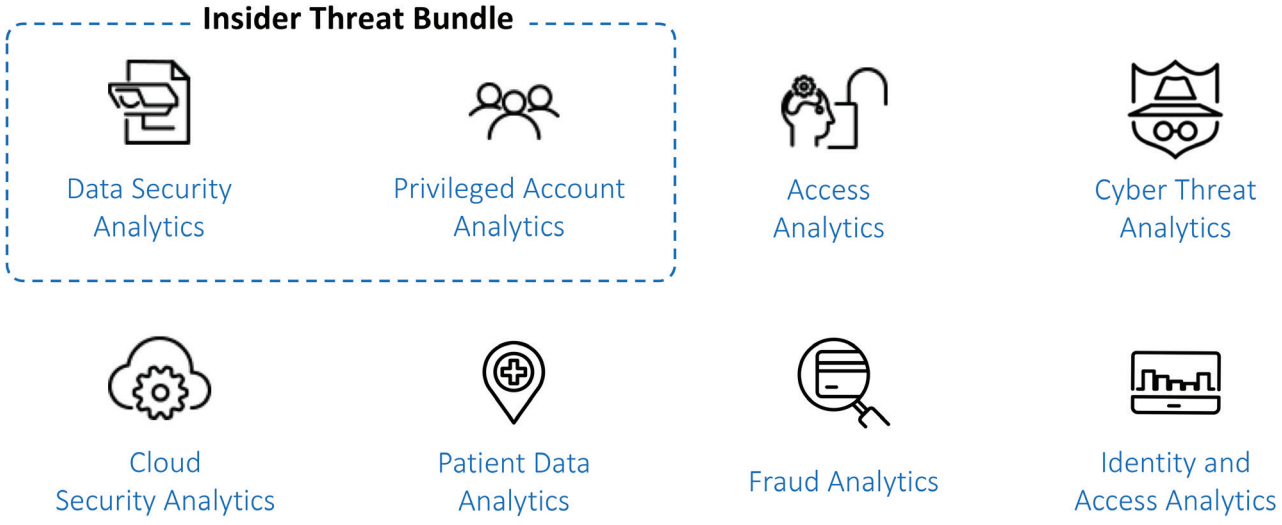

Figure 7: Securonix Packaged Apps

The figure above outlines some of the key packaged apps aligned to Securonix use cases. Securonix continues to invest in innovating and enhancing the content in order to ensure our customers always have the latest threat protection in place.

In Conclusion

Evolving threats require evolving methods to counter those threats. As malicious agents spread out their attacks in order to counter traditional, single-view security systems, an analytics-based approach is needed to counter them. Security analytics using threat chains with behavioral analytics at their core is the way forward for cybersecurity. Noise is reduced dramatically, from millions to single digits in some instances.

Without analytics, it is impossible to take on the new cybersecurity threat landscape. Reach out to us to find out more about how Securonix is revolutionizing the way enterprises perform cybersecurity, and about Securonix's comprehensive cybersecurity management platform.

ABOUT SECURONIX

Securonix is redefining the next generation of security monitoring using the power of machine learning and big data. Built on Hadoop, the Securonix solution provides unlimited scalability and log management, behavior analytics-based advanced threat detection, and intelligent incident response on a single platform. Globally, customers use Securonix to address their insider threat, cyber threat, cloud security, fraud, and application security monitoring requirements.

CONTACT SECURONIX

www.securonix.com
info@securonix.com | (310) 641-1000

0919


SECURONIX™