IANS

Security Intelligence: Can "Big Data" Analytics overcome our blind spots? An IANS Custom Report

December 2012

LogRhythm™

Contents

Introduction

The realm of information security in 2012 and beyond is vastly different and more complex than in previous years; organizations have significantly more intricate infrastructures while still supporting legacy applications and systems. We have staggering quantities of data (security included) to sort through and retain. Due to budget cuts and staffing freezes, many security teams have to do more with less time and support. Data breaches and major compromise scenarios dominate the news and security blogs alike. In most large organizations today, the primary tool for monitoring and responding within the environment is a Security Information and Event Management (SIEM) platform that aggregates and correlates data from a variety of data sources. These systems are notoriously complex, with numerous false positives to tune and widely varying capabilities from one vendor to the next. Are these products up to the task? Or are new capabilities needed?

Some organizations are wondering if they are already breached, and most feel they've got a handle on security, but do they? IANS conducted an independent survey of security leaders in larger organizations, and survey results are shown below (all charts depict number of respondents to each selection):

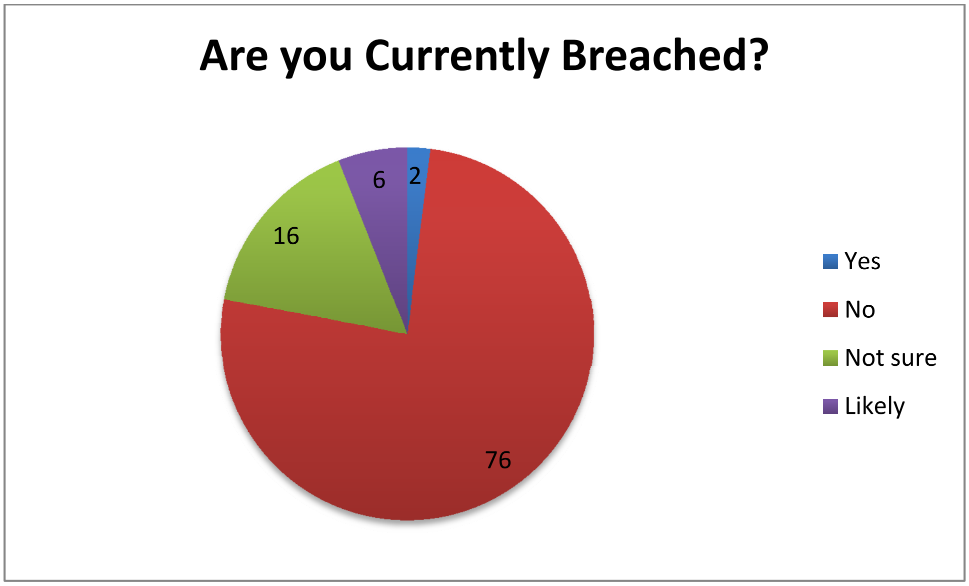

Couple these issues with the stark reality of the advancements in adversaries and threats facing us, and it paints a grim picture for security teams. Many IANS survey respondents felt that they were targets for advanced attackers in the last 12 months:

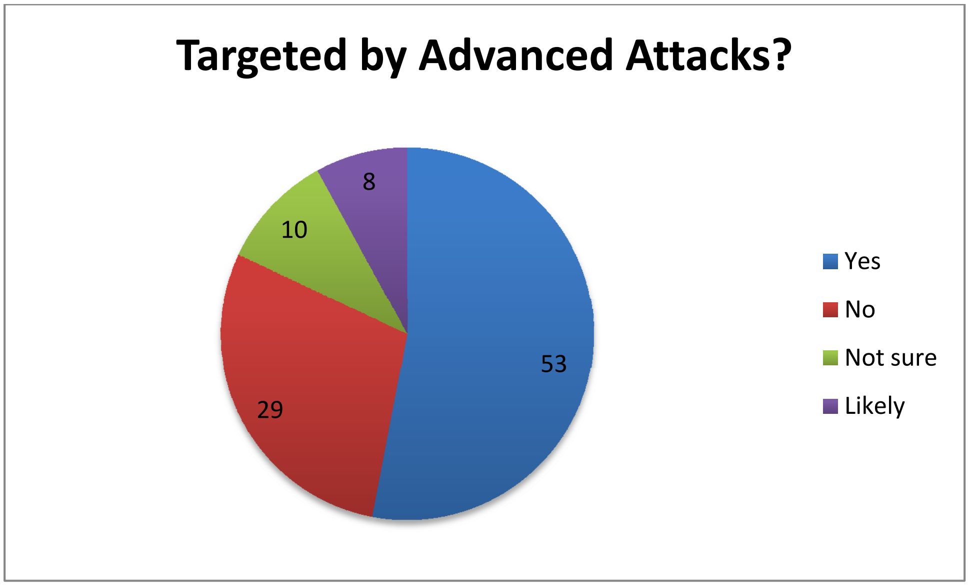

This data leads to some interesting questions. We know we're targeted, but feel strongly we haven't been breached. Is that a "head in the sand" perspective? Are most security teams really that confident that they aren't breached? **Or are they just afraid to publicly say so**? In this paper, we'll first explore the current threat and attack landscape, and then we'll look at some additional survey data in section titled "The State of Security Event Management and Monitoring" that may tell a very different story from the "overconfidence" shown here. We'll explore the current SIEM tools and capabilities we use to defend ourselves, and wrap up with some newer improvements in SIEM capabilities that might help turn the tide. Read on to find out the real story!

Today's Threat Landscape

In Verizon's annual Data Breach Investigations Report (DBIR) for 2012[1], a number of interesting statistics stand out. First, it is interesting that many organizations are initially being compromised by easy or only slightly difficult attacks. The majority of breach scenarios (65%) are accomplished through simple attacks, while another 24% are only moderately challenging for attackers. Even the subsequent actions are not particularly complex in most cases, with only 4% being classified as "high difficulty". However, we are still not doing a good job of preventing and/or rapidly detecting compromises. In addition, most attacks are now targeted against larger organizations (50%), with the number increasing to 63% for breached records.

While these numbers are interesting, there are many more conclusions to draw from the report. One in particular stands out - the timeframes from initial compromise to detection and detection to recovery. For all organizations, 54% took months to discover the initial breach, and larger organizations took months in 39% of cases and *years* in 9% of cases. Even more shocking is the method of detection; for all organizations, 92% were notified by an external party. This number improves somewhat for larger organizations, to 49%, but another 28% was detected passively inside the organization versus only 16% for active discovery efforts. Clearly we're not doing something right, but the fact is that the job is getting harder - fully 1/3 of Verizon's cases

[1] *2012 Data Breach Investigations Report (DBIR),* Verizon
http://www.verizonbusiness.com/about/events/2012dbir/

employed some sort of anti-forensics, such as encryption, timestamp modification, code obfuscation and modification, etc.

Nation state hacking is one of the significant trends affecting the threat landscape today. This is not a new trend, as we have major news stories of attacks originating from China and other nations as far back as 2007 (and even further). Many of these attacks are focusing on the power grid and utility sectors, as well as military and defense organizations ranging from companies to government agencies. In April 2009, the US electrical grid was compromised by Chinese and Russian hackers, and the US Joint Strike Fighter Program compromised through contractor networks. In 2010, the Stuxnet malware was discovered in Iran, affecting Siemens SCADA control systems. Another major target for these attacks seems to be the financial sector and other major technology companies in a variety of industries, however. Operation Aurora in 2010 was an attack against Google, Adobe, and other major technology firms by Chinese hackers who were seeking intellectual property and control of systems in these large firms.

Large data breaches continue to plague every industry and region in 2012, and will likely continue to do so for the foreseeable future. Several large and significant data breaches of note recently include:

- Sony Corporation was breached in 2011, leading to the compromise and exposure of 77 million Playstation Network accounts that included personal information and some financial data.
- Global Payments was compromised over a period of 6 months, leading to the loss of 1.5 million credit card numbers.
- Shanghai Roadway D&B Marketing Services Co. Ltd may have experienced an insider breach of 150 million customer records in December 2012.
- Gaming company Valve, Inc. was hacked, losing 35 million customer credit card numbers, passwords, and other personal information in 2011.

The prevalence of breaches and data points from sources like the Verizon report seem to indicate a major shift in attacker focus. Where we used to see attacks carried out for bragging rights and “hacker notoriety”, most attacks today seem focused primarily on financial and personal gain for attackers. Some seem to be more politically motivated, as in the large-scale attacks perpetrated by the Anonymous and LulzSec hacking groups over the past several years.

Advanced malware has played a huge role in the changing attacks experienced by more and more organizations. Even before Mandiant introduced the term “advanced persistent threat” into the information security lexicon, the security community has seen evidence of advanced malware tactics. Examples of advanced malware capabilities seen in the last several years include the following:

- Base64 encoded commands in Web pages used by botnets to download and control malicious files
- Self-destruct capabilities that remove all traces of the malware and leave no system modifications behind

- Insertion of encrypted Windows registry keys and multiple levels of version information
- Use of HTTP and HTTPS as transport protocols for command and control traffic
- Memory-resident attacks that do not alter operating system files and are resistant to forensics

Social engineering attacks such as phishing and pretexting have started to become more and more common as threat vectors as well, targeting end users directly. Spear phishing attacks that are highly focused and targeted (usually at high-level and executive users) have been prominent elements of some of the biggest breaches in the last several years, including the attack at RSA that led to compromise of their seed files for many customers' SecurID 2-factor tokens. These seed files were then used in a later attack against Lockheed Martin.

Insider threats are also a persistent challenge for many organizations today. While the percentage of external threat actors continues to grow steadily, according to the Verizon data breach report, insiders are often capable of much longer-term and stealthier fraud and compromise scenarios. Many "insiders" today are also unwilling participants due to social engineering and client-side attacks that turn their local laptop and desktop systems into vantage points for external attackers, as well, and this trend seems likely to continue.

The State of Security Event Management and Monitoring

SIEM systems are platforms that can take network, host, application, and security-specific data and analyze it with the goal of correlation for better understanding of trends, behaviors, and security events & incidents. Organizations that use SIEM and log management platforms have traditionally used them to aggregate large quantities of event data and create security-focused rules for analyzing and monitoring what's happening across the environment. Many larger organizations employ a Security Operations Center, or SOC, that employs SIEM tools as a primary mechanism for evaluating potential incidents and initiating investigations. Most organizations responding to the IANS survey felt that they had somewhat mature security event and incident monitoring and management programs in place:

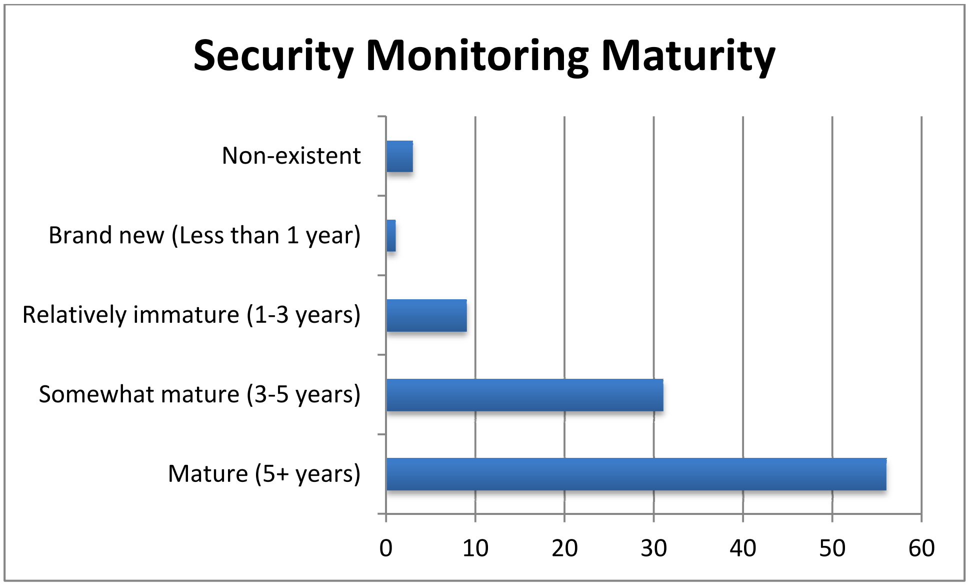

Note that most security teams feel that they have very mature monitoring programs already in place, and respondents felt that they were not currently breached despite the acknowledgement that advanced threats were targeting them (from the Introduction). So, does this mean that existing solutions are working? Read on!

SIEM components and focal areas

SIEM platforms invariably include the following architectural and process elements:

- Input sources for information analysis: These are likely to include IDS and IPS events, firewall and network device logs, network flow data, host system logs, application logs, anti-malware logs and alerts, vulnerability assessment results, and more. 71% of IANS survey respondents currently use flow data in their event monitoring strategies.
- Data normalization and storage: Normalization is the process of converting logs and events to a common format for analysis and interpretation. Some SIEM platforms include log management and data storage, as well, but many will use data from a larger central storage platform like a Storage Area Network (SAN).
- Data correlation and analysis: Security event monitoring platforms need to sift through normalized data and find patterns that trigger events.
- Alerting and response: Once correlation rules and monitoring thresholds have been triggered, events should create alerts that can initiate incident response efforts and additional monitoring.
- Forensics (varying degrees & types): With the advent of larger data sets and more advanced capabilities, most security event management platforms now allow for network and data forensics in one way or another. 74% of IANS survey respondents feel they can reconstruct events and create "after the fact" scenarios from event data.

- Reporting: Security event management and monitoring tools need to provide a broad array of reports ranging from the highly technical and granular (for security analysts) to more compliance-focused.

This will likely remain the same for newer platforms, although the variety and breadth of event sources will probably continue to expand.

Typical use cases and monitoring scenarios

Security event management platforms are evolving all the time, evolving from a standalone monitoring platform to one that provides true context awareness and analytics capabilities. Enrichment of event data (log data) with add-on systems such as Identity Management, Vulnerability Assessment, Configuration Management, and any other data sources that can add context to an event can help analysts detect events and respond faster. Simple examples include:

- Correlating DNS, WINS, and NIS Services to map IPs to names
- Correlating geo-location data to map IPs to geographical locations
- Correlating Active Directory or LDAP Services to map usernames to user identities
- Correlating vulnerability assessment information to map events with known vulnerabilities

Other common use cases and monitoring scenarios might include watching Windows and Unix event logs over time, monitoring network events from firewalls, IDS/IPS, and flow data, tracking user-specific event data like logon/logoff times and activities and objects accessed/modified, and matching vulnerability analysis data to attacks against specific systems and applications. IANS survey respondents indicated the following capabilities were currently within their team's grasp:

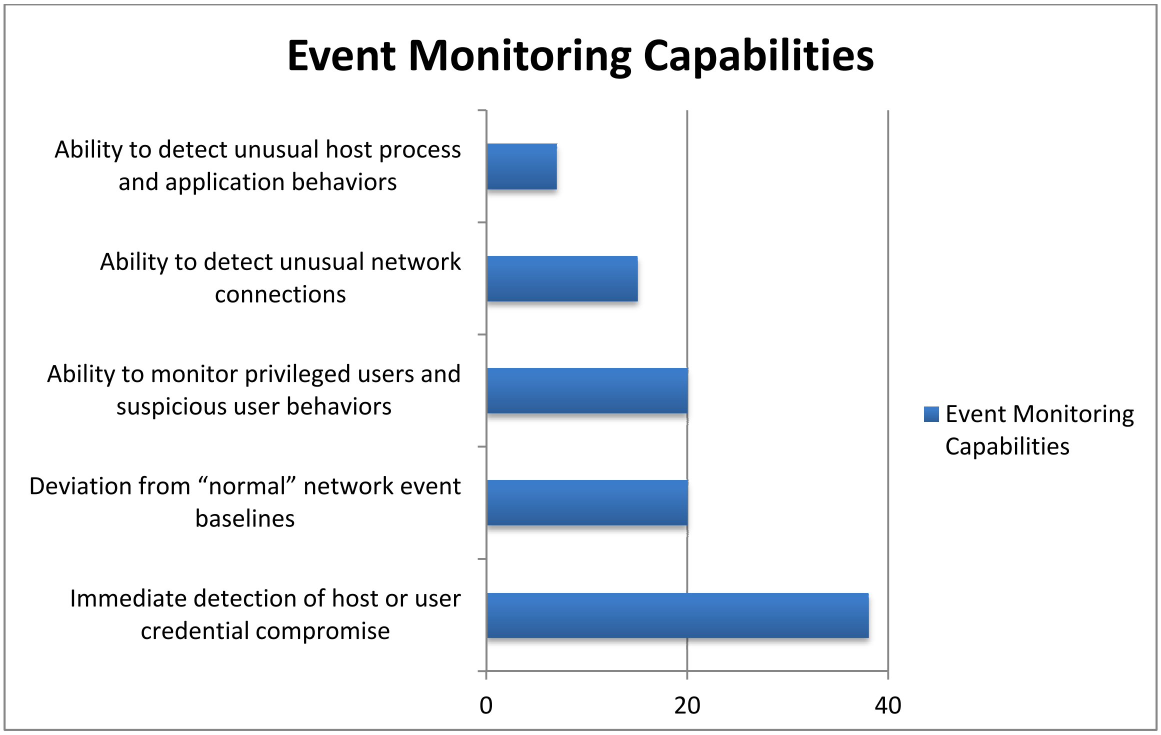

Here, finally, we have some data that really calls into question the previous answers supplied by respondents. Although many large organizations feel they have mature monitoring programs, and that they have avoided a breach thus far (from advanced attackers, as well), less than 40% feel they have the ability to immediately detect a compromised host or compromised user credentials, and even fewer feel they can detect unusual user and network behavior! These responses certainly lead us to feel that many organizations are relying on antiquated ideas of monitoring capabilities and maturity, especially given the severity of the types of malware and threat actors seen in the wild today.

IANS respondents also indicated specific that they were focusing on preventive and responsive controls, which demonstrates effort to implement better behavioral monitoring, anti-malware tools, and continuous monitoring of the entire network:

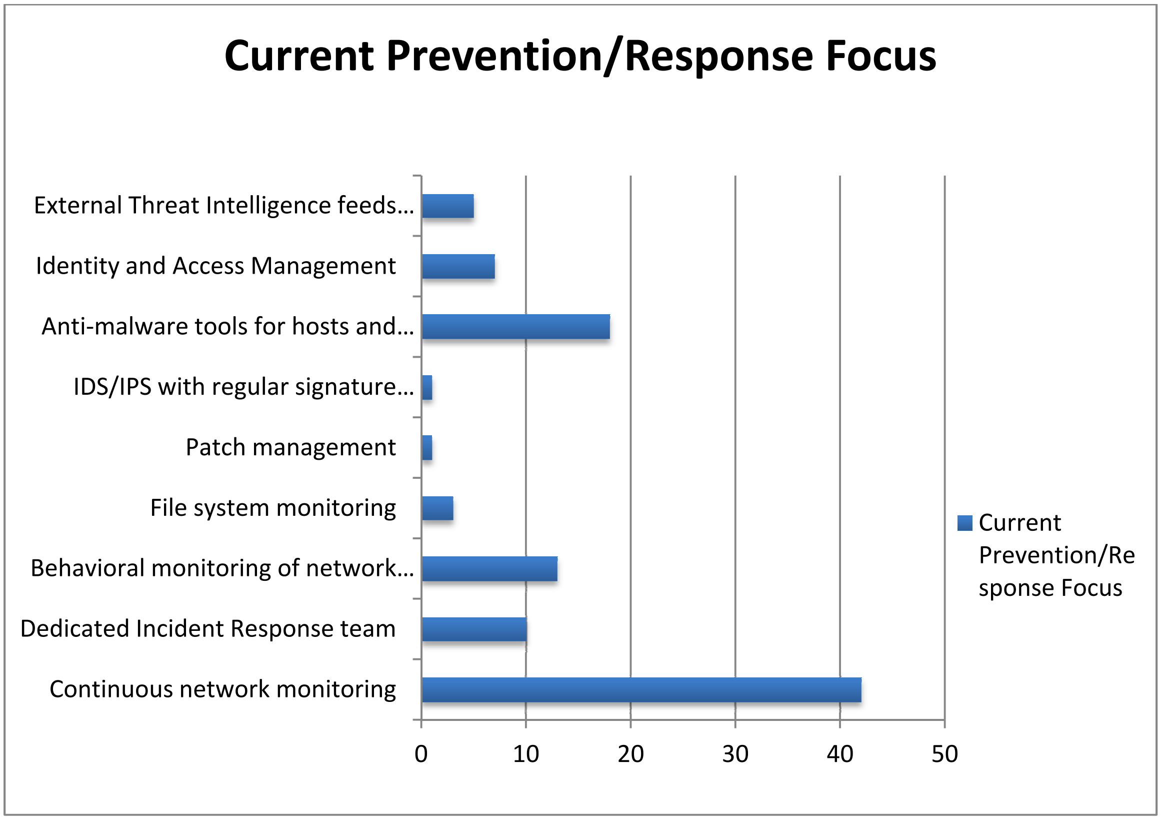

Blind Spots and Challenges

Each year, the SANS Institute publishes a Log Management survey[2] that also encompasses more robust security analysis tools like SIEM. In the 2012 report, survey respondents indicated that they considered identification of activities in “normal” background activity their biggest challenge, followed by correlation of multiple data types to meet complex monitoring needs and finally a lack of analytics capabilities:

[2] *SANS 8th Annual Log Management Survey*, SANS Institute
www.sans.org

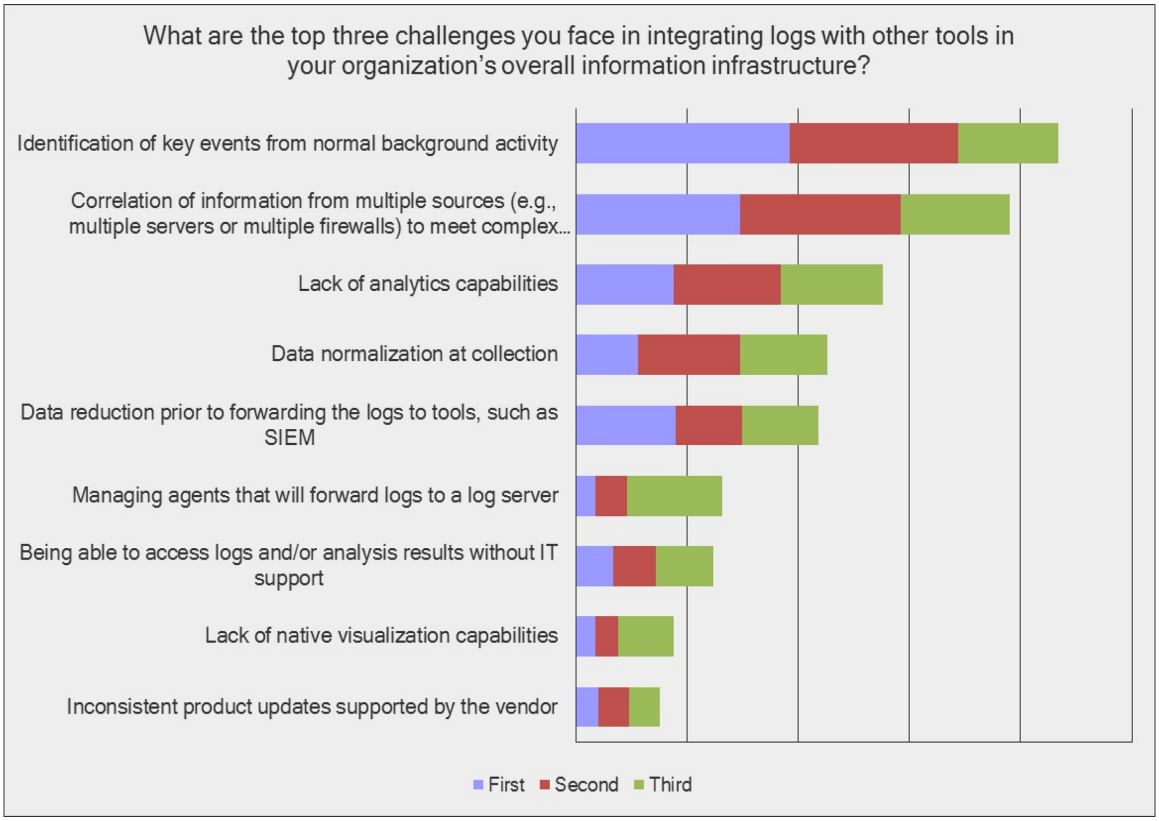

SANS 8th Annual Log Management Survey, SANS Institute, www.sans.org

Most organizations that responded to the IANS survey indicated that they were experiencing some significant challenges related to advanced threat detection and monitoring. Responses included the following:

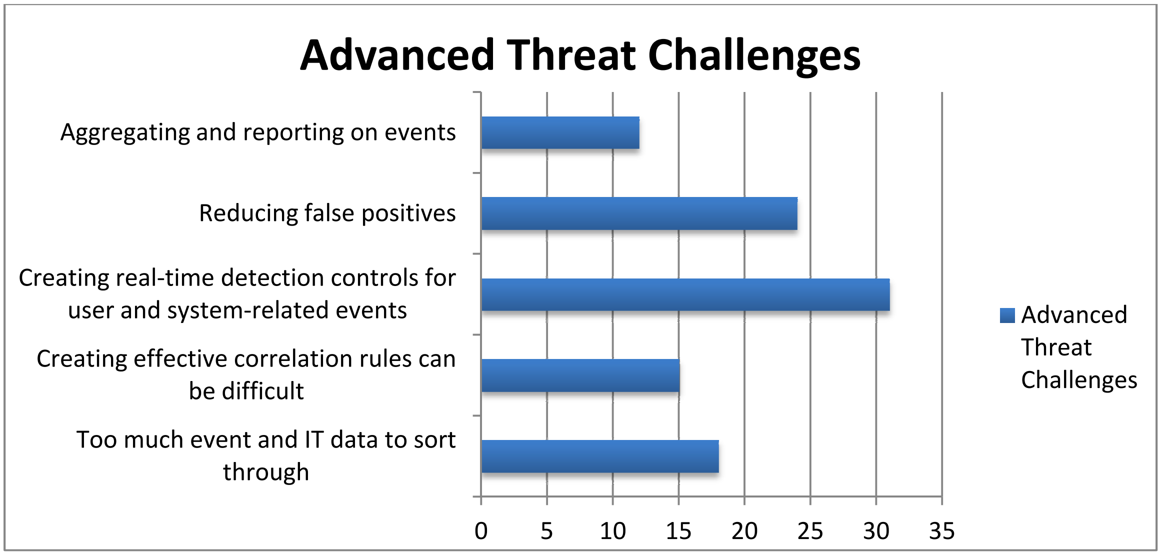

These results align with the data in the annual SANS study, showing that management of large data sets is a challenge, and that real-time controls for detection and false positive reduction are also major pain points. A majority of respondents to the IANS survey indicated that they were looking into security data analytics platforms that could assess large data sets and provide more comprehensive behavioral and baseline analysis:

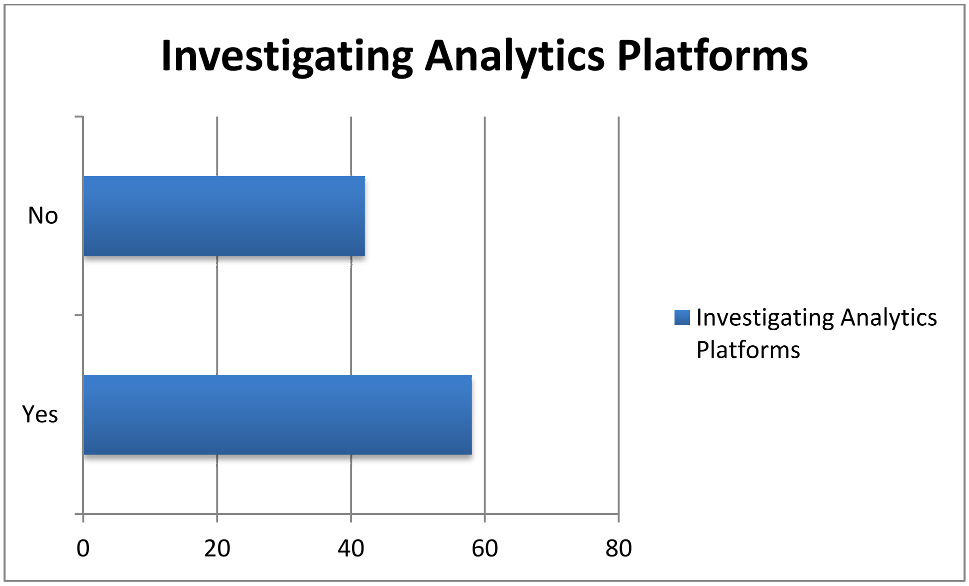

New Threats, New Security Needs

IT and security teams need tools that are just as advanced as the tools threat actors use. In the case of an event, intrusion analysts need to be able to quickly and accurately sort, normalize and analyze large sets of event information from multiple devices. These types of capabilities should also be available for alerting and forensics. There are a number of new capabilities that enterprise defenders need to help solve the problems illustrated in the IANS survey responses.

The first, and potentially biggest problem overall, is sorting through the "noise" and discovering trends and patterns indicative of potential incidents. Over time, security analysts have come to realize that signature-based tools can only identify so many things, and thus a major trend in security event management and monitoring technology is behavior baselining and analysis of event data. Another related need for analysis teams is false positive reduction. Most SIEM platforms are in constant need of tuning to eliminate the large numbers of false positives that continually manifest. In addition, even though some events may be interesting, they are often not truly indicative of an incident occurring or about to happen. One technique that can significantly help to reduce the overall quantity of alerts and lend more meaning to the alerts that do appear is behavioral whitelisting. Behavioral whitelisting allows analysts to go beyond simple network flow data and leverage other behavioral profile data along with event trends and attributes to determine what is "allowed", and start filtering all else quickly. For example, if trying to determine unusual behavior related to user account activity, analysts could include the following event data to more accurately pinpoint what "normal" is and then look for anomalies:

- Time of day for events
- Where events are originating from, and what patterns and typical profiles these sources fall into
- What roles or groups the user is a member of - is she a member of a more "sensitive" group like Finance or HR, or working in a call center?

- Critical data typically accessed by the user and groups/roles related to her, and what "normal" access patterns to that data look like, including frequency of access, from where access occurs, and specific users commonly accessing the data
- Thresholds defined for "normal" versus "abnormal" account activity or data access

With much more robust event correlation and analysis capabilities, behavioral profiling and whitelisting could prove to be invaluable in helping weed out events that aren't as useful in an investigation.

In regard to investigations, analysts definitely need more intuitive query and investigation tools at their disposal. While "ease of use" may seem like a trivial topic related to security event management, most analysts can tell you that SIEM platforms tend to be needlessly complex to use and manage. This can make rapid identification of events of interest and more in-depth investigation difficult and unwieldy, and as information security teams need to respond faster than ever before, this is less than ideal. An intuitive interface with rapid drill-down capabilities, sorting and graphing functions, and a variety of ways to analyze data sets and different types of events is invaluable.

Other trends on the horizon for security event management include more intuitive reporting and built-in "knowledge base" capabilities that can help defenders quickly and easily leverage the broad expertise of platform vendor engineers and the security community at large. A related topic that is becoming more and more relevant is the concept of defensive information sharing, where analysts from different organizations around the world can share correlation rules and event types of note with one another.

Conclusion

The threats facing our organizations are growing ever more complex. Based on the results from IANS independent survey, some of the following conclusions can be drawn about the current challenges teams face in combating advanced threats:

- Enterprise security teams are being targeted by advanced threat actors; in some cases, organizations feel they may already be compromised
- Large data sets, false positives, and rapid detection capabilities are some of the more pressing priorities in security event monitoring and management
- Most organizations are planning to look into security tools that can help them manage large volumes of security data and develop better security analytics

As the threats we face get more sophisticated, so too must the security tools we use to detect and combat them. More intuitive SIEM tools that combine behavioral analysis and whitelisting, "big data" analytics for both real-time threat/breach detection and after-the-fact forensic search/investigation, and more collaborative information sharing and knowledge creation for analysts will go a long way to helping us improve our defenses.

About IANS

IANS is the leading provider of in-depth security insights delivered through its research, community, and consulting offerings. Fueled by interactions among IANS Faculty and end users, IANS provides actionable advice to information security, risk management, and compliance executives. IANS powers better, faster, technical and managerial decisions through experience-driven advice.

IANS was founded in June of 2001 as the Institute for Applied Network Security. Inspired by the Harvard Business School experience of interactive discussions driving collective insights, IANS adapted that format to fit the needs of the information security professionals.

About LogRhythm

LogRhythm is the largest and fastest growing independent Security Information and Event Management (SIEM) provider in the world. The company's patented and award-winning SIEM 2.0 Big Data Security Analytics platform empowers organizations around the globe to detect breaches and the most sophisticated cyber threats of today, faster and with greater accuracy than ever before. LogRhythm also provides unparalleled compliance automation and assurance as well as operational intelligence to Global 2000 organizations, government agencies and mid-sized businesses worldwide.

Positioned as a Leader in Gartner's 2012 SIEM Magic Quadrant and listed as a "Champion" in Info-Tech Research's 2012 SIEM Landscape Report, LogRhythm also earned a perfect, 5-star rating and this year's exclusive "BEST BUY" in the SC Magazine SIEM Group Test. Additional awards have included Computing Security's Bench Tested Solution of the Year, SC Labs' "Recommended" 5-star designation twice, SC Magazine's Innovator of the Year Award, Readers Trust Award for "Best SIEM" solution and "BEST BUY" designation for Digital Forensics. LogRhythm is headquartered in Boulder, Colorado with operations in Canada, Europe and the Asia Pacific region.