Report

Needle in a Datastack Report

intel Security

Table of Contents

Big Data is changing the face of the global business landscape, with few technologies left untouched by the opportunities it presents—and information security is no different. With the increasing sophistication of attacks and growing regulatory pressures, the variety, volume, and analytic needs of security data have grown beyond the capabilities of traditional information management systems. The sheer volume of security-related data facing an organization these days can make identifying a threat like looking for a needle in a haystack. Yet collecting more data can also play a transformational role in information security. Organizations need to become smarter at harnessing and sifting through this wealth of information to protect themselves from the unrelenting threats they face every day.

With security breaches and fraud incidents continuing to make headline news and cybercrime activity booming, businesses must take more intelligent steps to address increasingly sophisticated security threats. Organizations are finding that they have a need for more intelligence-driven security and are beginning to initiate Big Data security programs.

Security professionals must harness the potential of Big Data to identify new trends, patterns, and threats to their organizations. Tools must be adopted to provide the visibility needed for greater levels of security intelligence. We believe it is time to embrace Big Data Security.

Organizations are failing to identify security breaches as they happen.

In this day of antivirus, endpoint security, and intrusion prevention, why should organizations seek to mine the data that their security systems collect? The reality is that just as security evolves, so too do the myriad of stealthy attack mechanisms devised by cybercrime perpetrators. These groups are developing malware that patiently probes from both inside and outside the network, blending in with normal activity and camouflaging insider abuse. Advanced targeted attacks (ATAs), such as advanced persistent threats (APTs), could potentially be idling within a network for weeks or months without triggering security systems—and it is this type of threat that can only be detected through careful and continuous data analysis. The bigger the organization, the bigger the data pool, and the harder and more time-consuming it is to detect anomalies. But it is the time factor that is critical here. Once activated, ATAs are limited only by the bandwidth of their victim's networks, potentially allowing terabytes of compromised information to flow out of the network in minutes.

For this report, UK technology market research firm Vanson Bourne interviewed 500 senior IT decision makers. The report investigates how well organizations are positioned to address the challenges of managing security in a world of huge amounts and varieties of data. The research highlights the scale of the challenge and the business imperative of spotting and managing anomalous and potentially dangerous activity amid the colossal amounts of data traffic. Specifically, the report reveals an alarming lack of appropriate security monitoring systems which leaves organizations vulnerable to cybercrime during the normal course of business—often for several days. This is further compounded by misplaced confidence in the robustness of current cyberdefenses and the increasing exposure to advanced threats. As the volume of security data grows bigger, IT and security professionals need to ensure that they are collaborating closely, as Big Data threatens to reveal a worrying dichotomy between the two.

Gartner contends that Big Data creates business value by enabling organizations to uncover previously unseen patterns and to develop sharper insights about their businesses and environments, including information security. In this report we shed light on some of the major trends, and provide guidance and best practices for security professionals as they look to navigate the Big Data Security landscape.

Inability to Identify Security Breaches

One of the major findings of the *Needle in a Datastack* report is the inability of the majority of organizations to identify security breaches and security risks as they happen.

While it is true that businesses are detecting breaches, finding out the length of time it takes depends on which part of the business you ask. In the Vanson Bourne study, 35% of IT decision makers stated that they could detect a data breach within minutes of its occurrence. The study also revealed that 22% said it often takes a full day to identify a breach, and 5% stated that it can take up to a week. This means that, on average, it takes 10 hours for an organization to recognize a security breach.

However, in its *2012 Data Breach Investigations Report*, Verizon suggests that of the large organizations that had experienced a compromise, not a single one was able to identify the threat within hours or minutes. In fact, more than 27% took several days, just 24% took weeks, and an alarming 39% took months between initial compromise and discovery—and 9% took a year or more. Needless to say, the sheer amount of data that could be harvested in this time is substantial and immeasurably valuable. It also appears that even if organizations were able to spot breaches within hours, 72% of attacks can compromise systems or networks in just seconds or minutes. Data exfiltration (transmitting data out of the network) takes just minutes or seconds in 46% of cases, according to Verizon.

Perhaps just as worrying, however, is the obvious disconnect between the reality of what systems are actually capable of and what IT decision makers believe them to be capable of. Given the volume of threats that organizations are trying to repel every day, the fact that a security breach can go unnoticed for more than a working day is a serious concern. Data loss, stolen intellectual property (IP), system downtime, compliance failure, damage to brand reputation, and erosion of customer trust are just a few of the serious implications of failing to spot threats in real time.

This situation is further exacerbated by the growth in use of mobile technologies by corporate employees and the even greater delays in spotting breaches in mobile endpoints. While 44% of respondents said breaches were spotted in minutes on mobile devices, a fifth of security professionals surveyed said it would take a day to spot a security risk; in fact, the average time to spot a mobile security breach increased to 14 hours.

Consider the havoc that a security breach could cause an enterprise within the course of a working day, given data download speeds, terabytes of commercially sensitive information, and IP that could be stolen and systems brought down, not to mention the negative impact on brand reputation and customer trust.

Loss of customer trust (62%) was the most acknowledged consequence of a security breach, followed by damage to the brand and corporate reputation (52%). Regulatory challenges (41%), financial loss through lost customers and fines (40%), and reduced employee trust in security (36%) also ranked highly on IT managers' list of concerns.

Yet it is not just the external threats that can lead to breaches. Data breaches can also come from within an organization due to malicious intent or carelessness. For example, the study found that it would take an average of 41 hours to spot a database administrator (DBA) abusing permissions. That is more than one work week. For example, in the San Francisco Department of Telecommunications and Information Services (DTIS), an IT administrator created a private administrative account on systems within the city's FiberWAN project. The administrator kept the password a secret, locking the organization out of its network setup for 10 days. It cost the department in excess of $1 million to reverse the disruption and ultimately resulted in a criminal conviction for the administrator, not to mention the loss of public services and reputation damage for this city agency.

- 22% of businesses need one day to identify a security breach.
- 5% of businesses need up to a week to identify a security breach.
- On average, it takes 10 hours for an organization to identify a security breach.
- Misplaced security confidence is putting organizations at risk.
- The second major finding from our study is that IT managers are overly confident about the level of security in their organizations.

These findings were supported by the recent Cybercrime and Security Survey Report published by the Australian Government and CERT Australia. While the majority of attacks reported by the 450 businesses represented were believed to come from external sources, an alarming 44% were believed to have originated from within organizations. This report found that more than half of the respondents viewed the attacks to be targeted at their organization, with motives that included illicit financial gain (15%), hactivism (9%), using the system for further attacks (9%), using the system for personal use (6%), attacks from foreign governments (5%), personal grievances (5%), and attacks from competitors (4%).

Respondents were also asked what factors they thought may have contributed to the incidents. The most highly-rated reason was the use of powerful automated attack tools (14%), followed by exploitation of unpatched or unprotected software vulnerabilities (11%), and exploitation of misconfigured operating systems, applications, or network devices (10%). This serves as a reminder that internally focused cybersecurity controls and measures are also critical.

While it is undoubtedly wise to secure the perimeter, most traditional technologies cannot handle internal threats. Threats like these can only be prevented through careful analysis of behaviors viewed against what is considered "normal" within an organization. Careful monitoring of processes may have prevented situations like this.

While more than half (58%) of security professionals said they experienced some type of security breach in the last year, 73% claimed they could assess their security status in real time. Organizations also responded with confidence in their ability to identify insider threat threats in real time (74%), perimeter threats (78%), zero-day malware (72%), and compliance controls (80%).

However, when the study drilled down further into these assertions, it was discovered that only 35% of businesses could actually detect security breaches within minutes. In fact, of those who said they had suffered a security breach in the last year, just a quarter (24%) said they recognized it within minutes, while the average time taken to detect an actual breach was a staggering 19 hours. Finally, when it came to actually finding the source of the actual breach, only 14% stated that they could do so in minutes, while 33% said it took a day and 16% said it took one week.

When organizations were asked which, if any, security information and event management (SIEM) solutions they had in place, only a fraction of responses were perceived as genuine SIEM tools. A large proportion of respondents believed that standard antivirus and database security systems provided them with the appropriate level of protection and real-time insights. It is true that these traditional tools can block or repel many attacks, but they lack the ability to detect threats from within. Advanced SIEM correlates event, threat, and risk data to accurately detect attacks in progress, serves as an investigation platform, and produces compliance reports resulting from activity monitoring.

The *Needle in a Datastack* survey indicates that many businesses have applied a "check-box" approach to security, believing that if they have a basic security infrastructure in place, it will be sufficient to protect them. But the threat landscape is evolving rapidly. Organizations must ensure that they have a coordinated and integrated defense across networks, devices, applications, databases, and servers to address escalating and increasingly sophisticated threats, which originate both internally and externally.

To get the visibility they need, security information from all points of vulnerability must be gathered and analyzed in real time in order to identify correlations and patterns that indicate attempts to breach defenses. Having this intelligence after breaches occur may not prevent the damaging consequences that could result.

Approximately 75% of IT decision makers claim they can assess their security status in real time, however:

- Only 35% said they could detect security breaches within minutes.
- Of those who said they had suffered a security breach in the last year, just 25% said they recognized the breach within minutes.
- The average time taken to detect an actual breach was 14 hours.

Organizations store approximately 11 to 15 terabytes of data a week, but 58% of firms store this data for just three months. There is no "one-size-fits-all" best practice, but organizations should be aware that advanced threats can occur over months or years by going under the radar of many blocking technologies. Not retaining the data impedes an organization's ability to find, understand, and eliminate these insidious threats.

Only 14% said they could identify the source of a breach within minutes, while 33% said it took a day, and 16% said it took one week. The *Needle in a Datastack* study found that organizations are storing approximately 11 to 15 terabytes of data per week. To put that in perspective and to highlight the wealth of information that is being left insufficiently guarded, 10 terabytes is the equivalent to the printed collection of the US Library of Congress. That's a lot of data to analyze and manage. What makes the situation even more concerning is that 58% of firms are storing this invaluable data for less than three months.

Advanced threats take all shapes and sizes, with some taking months to activate. According to the *McAfee® Threat Report: Fourth Quarter 2013*, the appearance of new APTs accelerated in the second half of 2013. These threats infiltrate an organization's defenses, staying dormant for months at a time. Then, when the organization least expects it, they strike, sending confidential information out or bringing malware and viruses into the organization before returning to a dormant state, only to launch another attack in the future. The *New York Times* was victim to this kind of attack, which persistently attacked the organization over a four-month period, infiltrating its computer systems and stealing the passwords of its reporters and other employees.

Organizations must retain their security data for longer and apply analytics to reveal patterns, trends, and correlations to spot and deal quickly with APTs. By using analytics, businesses can spot and block threats in real time, but long-term analysis of vast amounts of security information will also ensure that even dormant threats are found quickly.

Best Practices for the Age of Big Data Security

Today's threat landscape poses many challenges for organizations. Whether it is the volume of data, sophistication of threats, or the lack of real-time visibility, organizations are not harnessing the full potential of the Big Data Security opportunity. This will change soon as organizations begin experiencing true business value from unlocking the hidden insights of their security data. What are key best practices for real-time threat intelligence in the age of Big Data Security?

- **Collect all security information**—To achieve risk-based security intelligence, address APTs, and improve security monitoring, businesses need to store and analyze the right information. This goes way beyond log management. Without an automated approach and high-performance systems, this can be a real challenge. Deploying technologies that provide intelligent detection and automated collection will give organizations greater external threat and internal user context.
- **Synthesize actionable insights in real time**—The volume, velocity, and variety of information has pushed legacy SIEM systems to their limit. Now, with the pressing need to clearly identify complex attacks, organizations need advanced analytics that go beyond pattern matching to true risk-based analysis and modeling backed by a data management system that can keep up with complex real-time analytics.
- **Store and investigate long-term trends**—While real-time analysis of data is essential to derive security value from SIEM, organizations also need to be able to research long-term trends and patterns. Beyond just finding a "needle in a datastack," APT detection needs to be even more granular to find the right needle in a stack of needles.

- **Increase threat visibility**—To be effective, SIEM analysis has to go beyond an IP address and should provide security professionals with an understanding of the nature of the external system. Many SIEMs support threat feeds, but even more important is the breadth of the threat feed and the way it is used. Effective threat feed implementations use this data to perform a real-time reputation check, immediately alerting upon interaction with a known threat and pulling the reputation of the external source into the risk score.
- **Customize SIEM solutions**—Organizations that have an advanced, easy-to-use SIEM also must be able to customize their SIEM deployment based on risk which results in a stronger opportunity to detect APTs, insider abuse, and other hard-to-uncover attacks. At a minimum, that process requires having an understanding of which data is sensitive, which services are most critical, and who trusted users are that access these systems and services. A strong SIEM solution will have a risk-based engine where these parameters can easily be added to make risk prioritization meaningful.
- **Monitor and block**—Many organizations are frequently confused about monitoring versus blocking. Successful businesses understand what they can block and what they cannot and put a monitoring program in place to detect threats that can leverage available services, data, and resources. This is the mantra of "prevent what you can, monitor what you can't." At the heart of any strong security program is the protection of the confidentiality, availability, and integrity of assets. An effective SIEM will orchestrate this monitoring through collection of all security-relevant events, align it to context, and perform analytics to detect suspicious or malicious activity.
- **Create synergy between IT and security**—There needs to be greater understanding and cooperation between security and IT. Security and IT convergence is not at the stage it should be in most organizations, and IT departments often believe assets are protected when, in fact, they are not.

Today, businesses operate in the age of Big Data, and this applies just as much to keeping organizations secure as it does to connecting to their customers.

The advanced threats facing organizations today and in the future demand collection of more security data, analysis with a greater level of sophistication for real-time threat management, and retention of both current and historical data to enable long-term analysis of trends and patterns as a way of spotting dormant or insider risks.

Relying on legacy data management approaches and antivirus and database tools are insufficient for monitoring security breaches and do not keep organizations secure. Attacks are too constant and too sophisticated for SIEM to be disregarded. Moreover, SIEM is shifting from being seen as a compliance tool to an essential component of an organization's security architecture.

As organizations strive to counter increasingly hidden and complex threats, many are relying on a patchwork of security tools that they believe are up to the task but, in fact, leave numerous gaps across the infrastructure—from the network and devices to servers to databases. This approach is no longer viable. For real-time visibility and analytics, predictive insights, long-term modelling, and quicker, more effective threat response, organizations need an integrated, multilayered approach to security.

Big Data holds many answers, but only if an organization has the capability to harness the growing volume of security information. In deploying a SIEM solution to analyze this data, organizations can thwart advanced threats in real time and detect stealthy, dormant threats. Welcome to the age of Big Data Security.

Methodology

Technology market research firm Vanson Bourne interviewed 500 senior IT decision makers in January 2013, including 200 in the US and 100 each in the UK, Germany, and Australia.

McAfee. Part of Intel Security.
2821 Mission College Boulevard
Santa Clara, CA 95054
888 847 8766
www.intelsecurity.com

Intel and the Intel logo are registered trademarks of the Intel Corporation in the US and/or other countries. McAfee and the McAfee logo are registered trademarks or trademarks of McAfee, Inc. or its subsidiaries in the US and other countries. Other marks and brands may be claimed as the property of others. The product plans, specifications and descriptions herein are provided for information only and subject to change without notice, and are provided without warranty of any kind, express or implied. Copyright © 2014 McAfee, Inc. 61098rpt_needle-datastack_0514B_ETMG