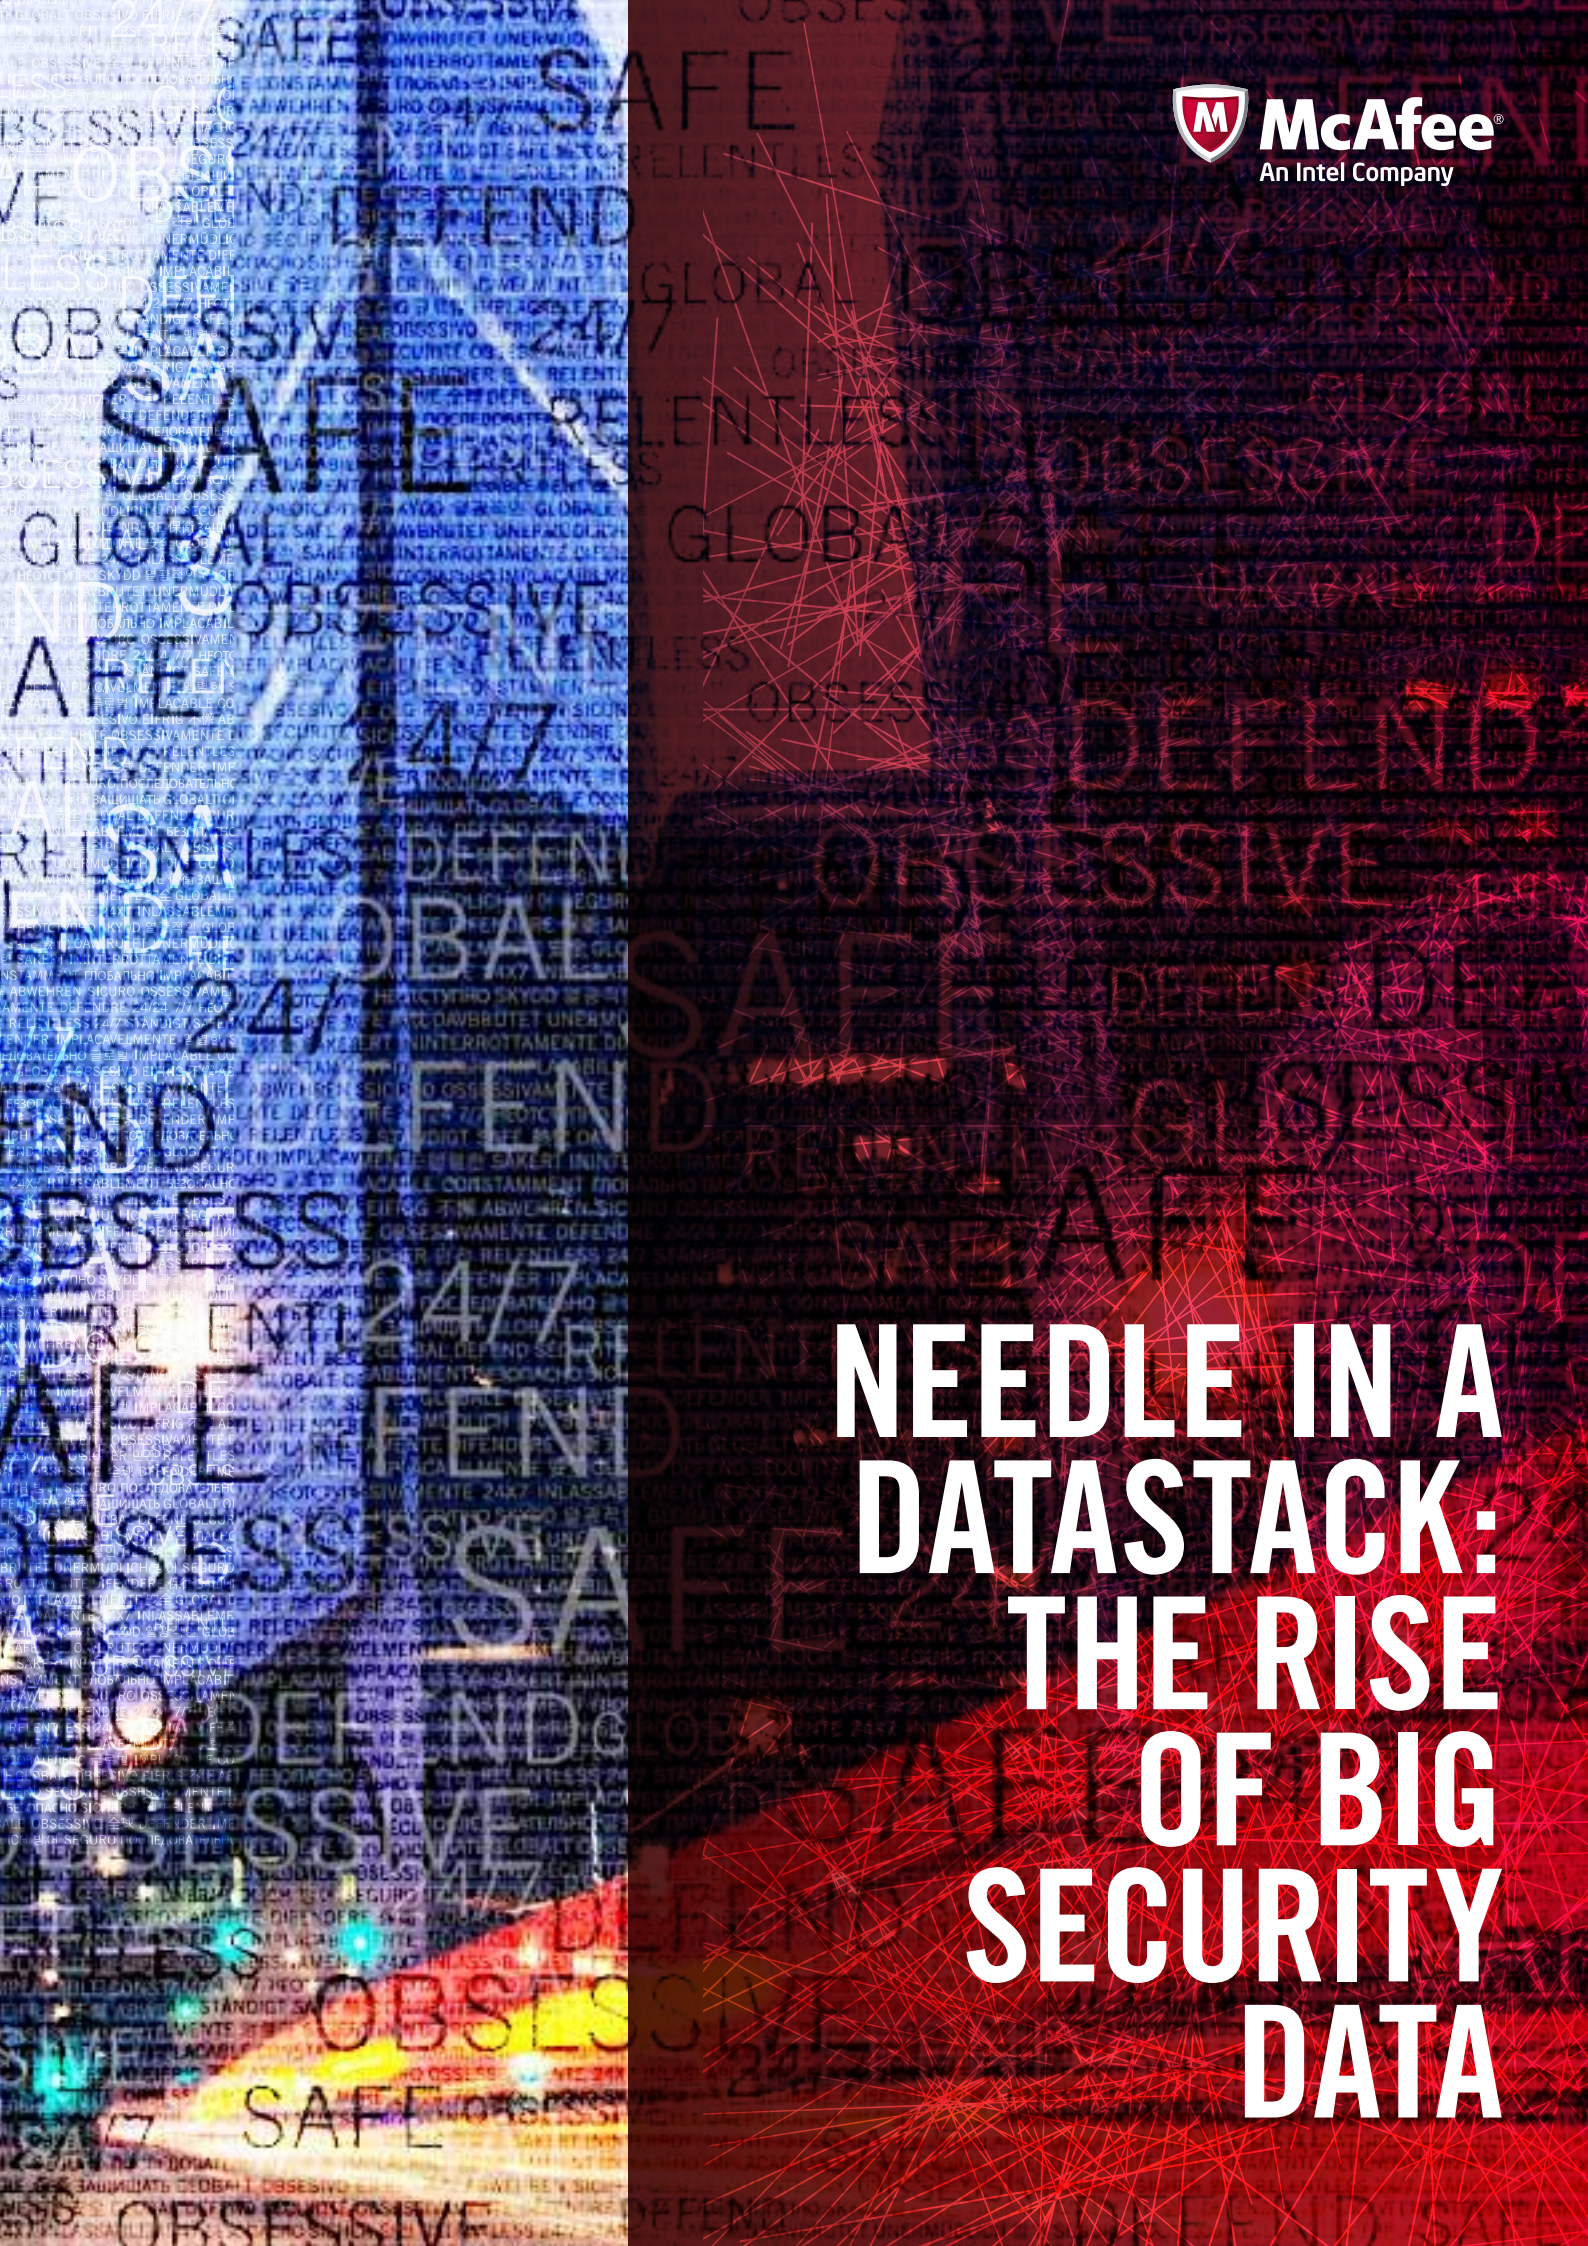
McAfee
An Intel Company
NEEDLE IN A DATASTACK: THE RISE OF BIG SECURITY DATA

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 ever-increasing sophistication of attacks and ever-growing regulatory pressure, the variety, volume and analytic needs of security data have grown beyond the capabilities of traditional information management systems. The sheer volume of security-relevant data facing an organisation 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 and organisations must become smarter at harnessing the right 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 the increasingly sophisticated security threats that face them. Organisations must embrace more intelligence-driven security and initiate big security data programmes.

In this day of virus protection, endpoint security and intrusion prevention, however, why should organisations seek to mine the data that their security systems collect? The reality is that just as security evolves, so too do the myriad, stealthy attack mechanisms devised by cybercrime perpetrators. These groups are prepared to play the long game, developing malware that patiently probes from both inside and outside the network, blending in with normal activity and camouflaging insider abuse. Advanced Targeted Attacks (ATA's) such as Advanced Persistent Threats could potentially lie 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 organisation, 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, ATA's are limited only by the bandwidth of their victim's network, potentially allowing terabytes of compromised information to flow out of the network in minutes.

500 interviews with senior IT decision makers were carried out for this report, Needle in a Datastack, which investigates how well organisations are positioned to address the challenges of managing security in a world of ever increasing amounts and types of data. Our research highlights the scale of the challenge but also 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 that leaves organisations vulnerable for entire working days to cybercrime. This is further compounded by a misplaced confidence in the robustness of cyber defences, and the increasing exposure to advanced threats. As volumes of security data grows ever bigger, IT and security professionals within organisations need to ensure that they are working closely together, as big data threatens to reveal a worrying dichotomy between the two.

Gartner contends that big data creates business value by enabling organisations 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 practice for security professionals as they look to navigate the big security data landscape.

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

ORGANISATIONS FAILING TO IDENTIFY SECURITY BREACHES AS THEY HAPPEN

One of the major findings from the Needle in a Datastack study was the inability of the majority of organisations to identify security breaches and security risks as they happen.

But while it is obvious that businesses are detecting breaches, finding out the length of time it takes depends on which part of the business you ask. In our study, 35 percent of IT decision makers stated that they can detect data breaches within minutes of them happening. This was further compounded by the revelation that more than a fifth (22 percent) said they would need a day to identify a breach, with a further 5 percent claiming it would take up to a week. This means that on average it takes 10 hours for an organisation to recognise a security breach.

However, in its 2012 Data Breach Investigations Report, Verizon suggests that of large organisations compromised, not a single one was able to identify the threat within hours or minutes and that in fact over a quarter (27 percent) took days, just under a quarter (24 percent) took weeks and an alarming 39 percent took months between initial compromise and discovery. Not to mention the 9 percent that 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 organisations were able to spot breaches within hours, 72 percent of attacks will take just seconds or minutes to compromise data, and data ex filtration, the action of transmitting data out of the network, again takes just seconds or minutes in 46 percent of cases found by Verizon. Putting it another way, the horse has bolted a long time before most organisations even notice the gate is open.

Perhaps just as worrying, however, is the obvious disconnect between the reality of what systems are actually capable of and what IT decisions makers believe them to be capable of – something quite different all together which is discussed later in this report.

Given the volume of threats that organisations are trying to repel every day, the fact 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 customer trust are just a few of the serious implications to a business from failing to spot threats in real-time.

What exacerbates this situation further is the growth in mobile technologies and the even greater delays found in spotting breaches through mobile endpoints. While more respondents (44 percent) said breaches were spotted in minutes via mobiles – probably due to the relatively small number of mobile endpoints and current working practices that most organisations have with regards to mobile working – a fifth of firms again said it would take a day to spot a security risk and 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 could be stolen and systems brought down, not to mention the negative impact on brand reputation and customer trust.

In fact, loss of customer trust (62 percent) was the most acknowledged consequence of a security breach by respondents, followed by damage to the brand and corporate reputation (52 percent).

Regulatory difficulties (41 percent), financial loss through lost customer and fines (40 percent), and loss of employee trust in security (36 percent) also ranked highly on IT managers' list of concerns.

Yet, it is not just the external threats posed. Data breaches can come from inside an organisation, whether malicious or inadvertent. For example, the study found that it would take an average of 41 hours to spot a database administrator (DBA) abusing their permissions. That is more than one working week. Illustrating this threat is what happened in the City of San Francisco Department of Telecommunications and Information Services (DTIS), where Terry Childs, an IT administrator, created a private administrative account on systems within the city's FiberWAN project. Childs kept the password a secret, locking the organisation out

of its networking set-up for 10 days. It cost the department in excess of $1 million to reverse the disruption and ultimately resulted in a criminal conviction for Childs, not to mention the loss of public services and reputational damage.

These findings were supported by the recent CyberCrime & 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 originate from within organisation. This same report found that more than half the respondents viewed the attacks to be targeted at their organisation – with motives being illicit financial gain (15%), hactivism (9%), using the system for further attacks (9%), using the system for personal use (6%), being from a foreign government (5%), personal grievance (5%), and being a competitor (4%). Respondents were also asked what factors they thought may have contributed to the incidents.
The highest 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%) and serves as a reminder that internally-focused cyber security controls and measures are also important.

While it is undoubtedly wise to secure the perimeter, no traditional technology can account for an internal threat – these can only stand a chance of being prevented through careful analyse of actions viewed against what is considered "normal" within an organisation. Careful monitoring of processes may have prevented situations like this, stopping suspicious activity before it was too late.

22 percent of businesses need one day to identify a security breach

5 percent of businesses need up to a week to identify a security breach

On average it takes 10 hours for an organisation to identify a security breach

MISPLACED SECURITY CONFIDENCE PUTTING ORGANISATIONS AT RISK

The second major finding from the Needle in a Datastack study is the misplaced confidence IT managers have in levels of security.

While over half (58 percent) said they experienced some type of security breach in the last year, nearly three quarters (73 percent) claimed they can assess their security status in real-time. Organisations also responded with confidence in their ability to identify in real-time insider threat detection (74 percent), perimeter threats (78 percent), zero day malware (72 percent) and compliance controls (80 percent).

However, when the study drilled down further into these assertions, it found that only 35 percent of businesses could actually detect security breaches within minutes. In fact, of those that said they had suffered a security breach in the last year, just a quarter (24 percent) had recognised 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 percent could do so in minutes, while 33 percent said it took a day and 16 percent a week.

Interestingly, when organisations were asked which, if any, security information and event management (SIEM) solutions they had in place, only a fraction of responses were perceived as genuinely SIEM tools. A large proportion of organisations 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 provide the blocking power to repel many attacks, but critically they also lack the monitoring power capable of identifying the threat from within.

Security information and event management (SIEM) brings event, threat, and risk data together to detect attacks in progress, serve as an investigation platform, and produce compliance reports related to activity monitoring.

The Needle in a Datastack study indicates that many businesses have applied a ‘tick-box’ approach to security, believing that if they have a basic security environment that will be sufficient to protect them. But the threat landscape is evolving rapidly. Organisations must ensure that they have a coordinated and integrated defence across networks, devices, applications, databases and servers to address the broad, escalating and increasingly sophisticated threat landscape, both internally and externally.

To have the visibility required, security information from all points of vulnerability must be gathered and analysed in real-time to identify correlations and patterns that indicate attempts to breach defences. Having this intelligence after the event will be too late to prevent the damaging commercial consequence that could result.

Three quarters of IT decision makers claim they can assess their security status in real-time but:

- **Only 35 percent of businesses could detect security breaches within minutes**
- **Of those that said they had suffered a security breach in the last year, just a quarter recognised it within minutes**
- **The average time taken to detect an actual breach was 14 hours**
- **Only 14 percent could identify the source of a breach within minutes, while 33 percent said it took a day and 16 percent a week**

The Needle in a Datastack study found that organisations are storing approximately 11-15 terabytes of data a week. To put that in perspective and to highlight the wealth of information that is being left insufficiently guarded, 10 terabytes is the equivalent of the printed collection of the Library of Congress. That's a lot of data to analyse and manage. What makes the situation even more concerning is that 58 percent of firms are storing this invaluable data for less than three months.

Highly developed threats take all shapes and sizes, with some taking months to activate. According to the McAfee Threats Report: Fourth Quarter 2013, the appearance of new Advanced Persistent Threats (APTs) accelerated in the second half of 2012. These threats infiltrate an organisation's defences, undetected for months at a time, sitting dormant. Then when the organisation least expects it they strike, sending confidential information out or bringing additional malware and viruses into the organisation before returning, until the next time, to a dormant state. For example, the New York Times was recently victim to this kind of attack, which persistently attacked the organisation over a four month period, infiltrating its computer systems and stealing the passwords of its reporters and other employees.

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

Organisations store approximately 11-15 terabytes of data a week but 58 percent of firms store this data for just three months

There is no "one-size-fits-all" best practice but organisations should be aware that advanced threats can occur over months or years, going under the radar of many blocking technologies – not retaining the data eliminates the ability to find them

BUSINESSES INCREASINGLY EXPOSED TO ADVANCED PERSISTENT THREATS

BEST PRACTICE FOR THE BIG DATA SECURITY AGE

Today's advanced threat landscape poses many challenges for organisations. Whether it is the volume and sophistication of threats across all fronts or the lack of real-time visibility, security appears to be blinkered to the big data opportunity. But this must change, and quickly, as organisations need to turn their attention to harnessing and unlocking the hidden insights of their security data.

So what are the 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 advanced persistent threats and improve security monitoring, businesses need to store and analyse the right information. This goes way beyond log management. Without an automated approach and high-performance systems, this is a real challenge. Deploying technologies that provide intelligent detection and automated collection will give organisations 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, organisations need advanced analytics that go beyond pattern matching to true risk-based analysis and modelling that is backed by a data management system that can create complex real-time analytics.

Store and investigate long term trends – while real-time analysis of data is essential to derive security value from SIEM, organisations also need to be able to research long-term trends and patterns. Beyond just finding a 'needle in a datastack', APT detection can go even more granular to find the right needle in a stack of needles.

Threat visibility – to be effective, SIEM analysis has to go beyond an IP address and understand the nature of the external system. While many SIEMs support threat feeds, the breadth of the threat feed and the way it is used is important. Effective threat feed implementations use this data to perform a real-time reputation check – immediately alerting on interaction with a known threat and pull the reputation of the external source into the risk score.

Customisation – organisations that take a few extra steps to customise their SIEM deployment based on risk have a stronger opportunity to detect APTs, insider abuse and other hard to find attacks. At a minimum that process requires having an understanding of what data is sensitive, what services are most critical, and who their trusted users are with access to these systems and services. A strong SIEM solution will have a risk-based engine where these parameters can easily be added to make risk prioritisation meaningful.

Monitor and block – many organisations are frequently confused between the ability to monitor and to block. Successful businesses understand what they can block – and what they can't – and put a monitoring programme in place to detect threats that can leverage available services, data and resources. This is effectively the mantra of 'prevent what you can, monitor what you can't'. At the heart of any string security programme is the protection of the confidentiality, availability and integrity of assets. An effective SIEM will orchestrate this monitoring through collecting all security relevant events, align it to context and perform analytics to detect suspicious or malicious activity.

IT/Security synergy – 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 organisations and the IT department often believes assets are protected when actually they're not.

Businesses operate today in the age of big data and this applies just as much to keeping organisations secure as it does to connecting to their customers.

The advanced threats facing organisations today and in the future demand collecting more security data, analysing it with a greater level of sophistication for real-time threat management, and importantly keeping the data longer to enable long-term analysis of trends and patterns to spot dormant or insider risks.

Legacy data management approaches, and believing that antivirus and database tools are sufficient for monitoring security breaches, will not keep organisations safe. Attacks are too constant and too sophisticated for SIEM to be ignored as a layer of every organisation's defence. Moreover, SIEM is shifting from being seen as a compliance tool to an essential requirement that plays an everyday, critical role in the security architecture of an organisation.

As organisations battle to counter increasingly 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 of vulnerability across the infrastructure from the network and devices to servers and databases. This patchwork approach is no longer sufficient. For real-time visibility and analytics, predictive insights, and long-term modelling, organisations need an integrated, multi-layered approach to security.

Big data holds many answers, but only if an organisation has the capability to harness the ever growing amounts of security information. In deploying a SIEM solution to analyse this data, organisations can repel advanced threats in real-time as well as spot the stealthy, dormant threat. Welcome to the age of big security data.

Methodology
Vanson Bourne interviewed 500 senior IT decision makers in January 2013, including 200 in the USA and 100 each in the UK, Germany and Australia.

BIG DATA HOLDS MANY ANSWERS, BUT ONLY IF AN ORGANISATION HAS THE CAPABILITY TO HARNESS THE EVER GROWING AMOUNTS OF SECURITY INFORMATION. IN DEPLOYING A SIEM SOLUTION TO ANALYSE THIS DATA, ORGANISATIONS CAN REPEL ADVANCED THREATS IN REAL-TIME AS WELL AS SPOT THE STEALTHY, DORMANT THREAT.

WELCOME TO THE AGE OF BIG SECURITY DATA.

McAfee
An Intel Company

Needle in a Datastack: The rise of big security data