When Minutes Count

How experts fight advanced threats with real-time SIEM and identification of eight key indicators of attack.
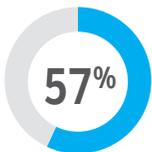
(intel) Security

Contents

Time is crucial.

78%
Of those who can detect attacks within minutes **use a real-time SIEM solution.**

57%
Could detect within minutes and had 10 or fewer targeted attacks.

12%
Of agile organizations investigated more than 50 incidents last year.

52%
Of companies that are least concerned about attacks have a real-time SIEM.

5 out of 8
most useful indicators of attack **rely on time.**

Executive Summary

The steady parade of attacks disclosed in 2014 has made it abundantly clear that current security efforts and checkbox compliance aren't sufficient for data protection. In a new survey commissioned by McAfee, part of Intel Security, 74% of respondents remain highly concerned about their ability to handle targeted attacks and advanced persistent threats (APTs). Over half of respondents had investigated 11 or more targeted attacks in the previous year.

Despite the furor, or perhaps because of it, it's not always easy to tell where to invest for results and peace of mind. To turn fear, uncertainty, and doubt into a prescription for proactive attack prevention, McAfee commissioned Evalueserve to perform a health check on organizational abilities to deal with advanced and targeted attacks. We also collected best practice guidance based on incident response data and penetration tests performed at enterprise sites by Foundstone Professional Services experts.

Time is a crucial success factor:

- 78% of those able to detect targeted attacks within minutes use a real-time security information and event management (SIEM) solution.
- 57% of companies that could detect targeted attacks within minutes experienced 10 or fewer targeted attacks last year.
- 12% of agile organizations (those detecting attacks in minutes) investigated more than 50 incidents last year, underscoring the effort being expended by both attackers and defenders.
- 52% of those least concerned about attacks have a real-time SIEM.
- 5 of the 8 most useful indicators of attack (according to Foundstone investigations) rely on time as a meaningful attribute of an event.

We found that the most effective organizations focused on several key indicators to detect attacks:

- Unusual alert patterns can help organizations detect reconnaissance, weaponized malware, compromised assets, and remote control activities.
- Suspicious outbound traffic shows compromised hosts, command and control, and exfiltration.
- Unexpected internal traffic reveals stolen privileges, lateral movement, and propagation.

The final finding

Our survey indicates that companies with early attack detection skills are faring best against targeted attacks. Existing technologies—often existing deployments—are capable of delivering better protection and faster incident response than companies are achieving today. Many generate alerts and insights that could prevent a data breach or service disruption if the signal could be isolated from the noise. Most organizations just aren't using the available intelligence and tools to their fullest or constructing indicators of attack into a cogent picture in a timely fashion.

Evalueserve Survey Results

Pressures are high

In August 2014, a broad global survey of IT and security teams revealed that organizations large and small are struggling to build an effective strategy against targeted attacks and remain very worried about this threat area.

We wanted to understand if this worry was based on individual experience or impressions derived from the media. The data said the threat is real. A clear majority (58%) of companies with more than 50 employees said that they had investigated more than 10 targeted attacks in the previous year. This volume of investigations requires a significant resource commitment.
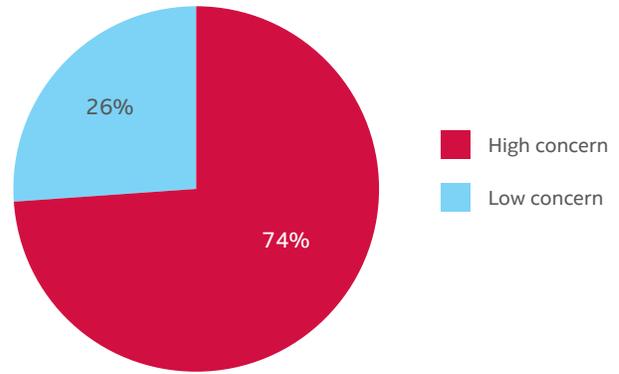

Figure 1. Evalueserve Question: Please rate how concerned you are about targeted attacks and APTs (looking at the past year in light of high profile breaches)?
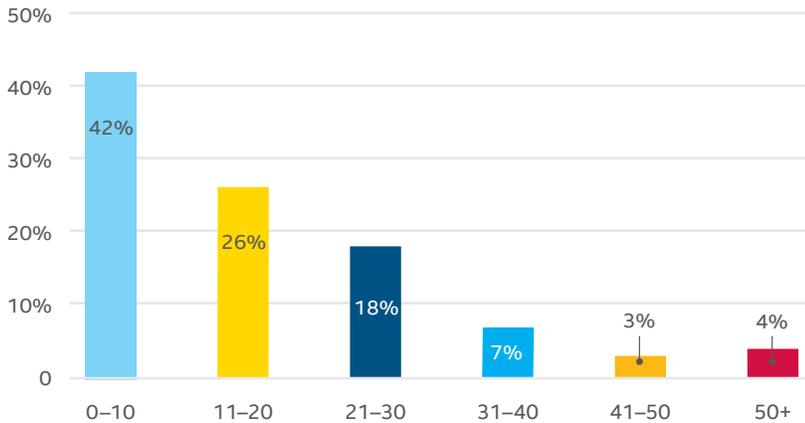

Figure 2. Evalueserve Question: How many times in the past year has your company investigated a suspected targeted attack or APT?

This concern makes sense. Exposure of regulated data has forced disclosure of many attacks in 2014, and those attacks have hit a broad swath of companies—from mom-and-pop shops to global brands—driving the hype. Two things happen as the obvious targets get better at defending themselves. Some determined actors work harder (12% of organizations detecting attacks within minutes investigated more than 50 incidents last year). Other criminals turn to smaller, less well-defended entities. This pattern was evident with spoofing of banking websites, for example. Once big national banks became poor targets, regional banks and credit unions fell prey. Phishing is following the same trajectory.[1]

However, it's not just about what hits the headlines. A SANS Institute study on incident response documented that 32% of companies affected by data breaches in the last two years had lost intellectual property (IP), a source of competitive advantage and profit margins.[2] These losses seldom see news coverage unless there is litigation, since regulations do not require disclosure of IP loss.

SIEM is a security intelligence solution designed to analyze security event, flow, and log data in real time for internal and external threat management, and also to collect, store, analyze, and report on log data for regulatory compliance and forensics.

Methodology

During August 2014, Evalueserve surveyed 473 IT decision makers from companies larger than 50 employees in the US, UK, Germany, France, and Australia. Online interviews captured awareness of advanced persistent threats, the frequency of threats witnessed and investigated, experiences with SIEM, and security practices.

Are we getting better? Detection speed is a metric for effectiveness.

It's clear security organizations are worried. We wondered if this angst was translating into action, if it meant people were getting faster, more effective, and more confident at detecting attacks. The Evalueserve survey investigated how organizations assess their own capabilities to protect or defend against these events. The good news is that 53% of organizations surveyed indicated discovery time of hours or minutes.
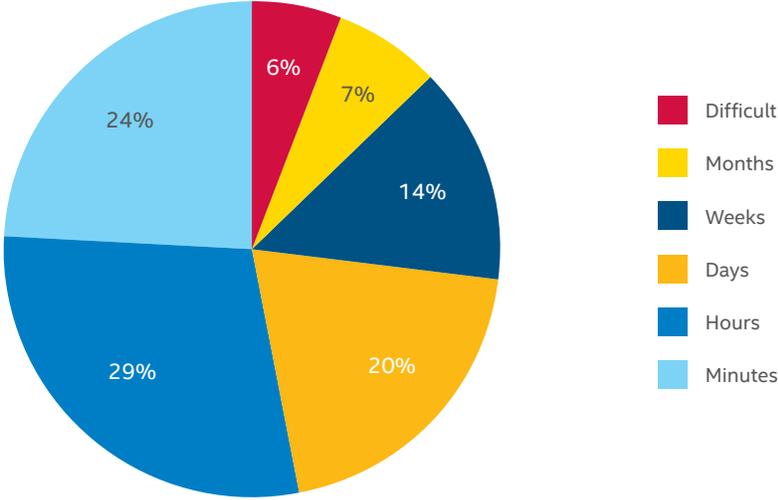

Figure 3. Evalueserve Question: Almost half of respondents take days or longer to detect an attack.

When we looked inside these numbers for patterns, one clear distinction for effectiveness was the presence of advanced SIEM technology. Although 93% of our respondents in companies larger than 50 had a SIEM, only half considered themselves to have an "adequate real-time proactive SIEM."
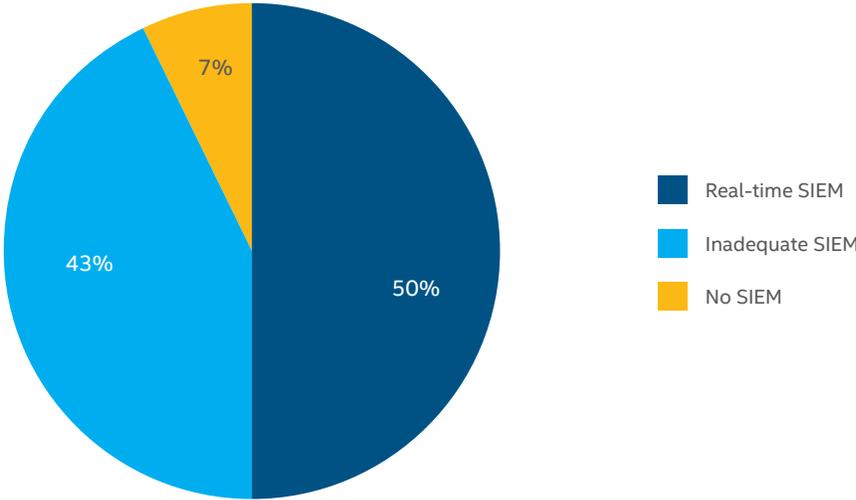

Figure 4. Evalueserve Question: Half of the organizations surveyed have no SIEM or an inadequate SIEM solution.

Success Story: Volusion

Lance Wright, information security manager for Volusion, cites a recent example of how his next-generation SIEM has changed the game for his team.

"We noticed a workstation making odd authentication requests to the domain controller at two o'clock in the morning. That could be normal activity, but it could also be a sign of something malicious. So we set up a rule to alert us if any workstation has more than five authentication requests during non-business hours to help us identify the attack early, before any data is compromised."[3]

Rather than wait for workstation compromise, Wright's team are actively scanning for the first hints of anomalous activity—they are now the hunters.

This demarcation can be explained since the original adoption of SIEM technology was for log archival for compliance. The emphasis was on passing audits, not mining the data at high speed. In contrast, modern SIEM solutions integrate threat intelligence, correlation, analytics, active response, and adaptive technologies that are specifically geared to help incident response.

Just as antivirus isn't sufficient for prevention, neither is having a SIEM active at its default settings or configured to merely store or mine historical data. Instead, real-time SIEMs enable rapid, complex manipulation to turn pure data into early attack detection and instant action on security events:

- Behavioral-based (rule-less) correlation can trigger priority alerts and automated responses based on risk scores tied to specific services and combinations of events, or thresholds of changes in these indicators.
- Baseline-driven anomaly detection fires based on atypical actions—once "normal" is defined, "abnormal" events can be given heightened visibility.
- Inclusion of external threat feeds enhances the internally sourced behavioral and baseline detection methods.
- Threat prioritization allows systems to score and initiate responses based on suspicious activities and the relevance of threats to specific assets due to asset value, vulnerabilities, patching levels, and countermeasures in place.

What difference does it make? Huge. 78% of those able to detect attacks in minutes had a real-time, proactive SIEM.
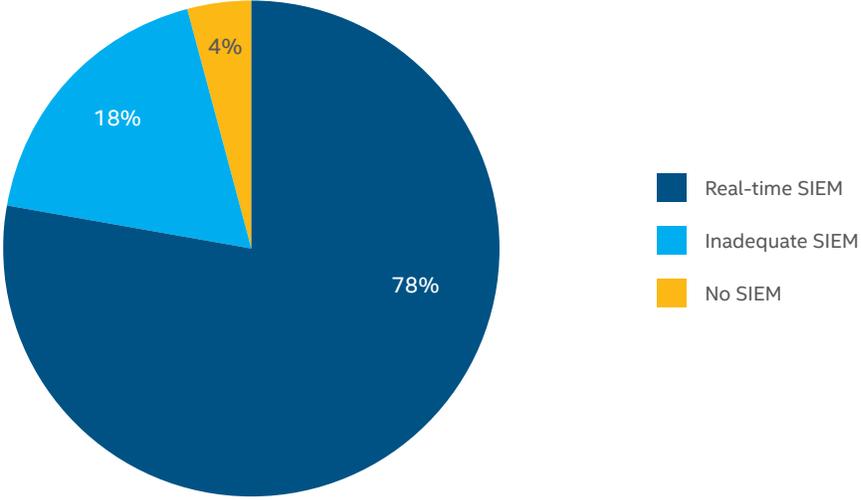

Real-time SIEM
Inadequate SIEM
No SIEM

Figure 5. Evalueserve Question: Use of SIEM by those detecting attacks in minutes.

Advice from the Front Line

The Evalueserve survey quantified individual practitioners' self-assessments and experiences. For an outside view, we turned to the experts. So far in 2014, Foundstone Professional Services consultants have assisted with more than 200 investigations, including several of the most publicized attacks.

We asked about indicators of attack (IoAs), events that could reveal an active attack before indicators of compromise become visible. Use of IoAs provides a way to shift from reactive cleanup/recovery to a proactive mode, where attackers are disrupted and blocked before they achieve their goal of data theft.

Eight indicators of attack

The following most common attack activities could have been used, individually or in combination, to diagnose an active attack:

1. Internal hosts communicating with known bad destinations or to a foreign country where you don't conduct business.

2. Internal hosts communicating to external hosts using non-standard ports or protocol/port mismatches, such as sending command shells (SSH) rather than HTTP traffic over port 80, the default web port.

3. Publically accessible or demilitarized zone (DMZ) hosts communicating to internal hosts. This allows leapfrogging from the outside to the inside and back, permitting data exfiltration and remote access to assets. It neutralizes the value of the DMZ.

> 4. Off-hour malware detection. Alerts that occur outside standard business operating hours (at night or on weekends) could signal a compromised host.

> 5. Network scans by internal hosts communicating with multiple hosts in a short time frame, which could reveal an attacker moving laterally within the network. Perimeter network defenses, such as firewall and IPS, are seldom configured to monitor traffic on the internal network (but could be).

> 6. Multiple alarm events from a single host or duplicate events across multiple machines in the same subnet over a 24-hour period, such as repeated authentication failures.

> 7. After being cleaned, a system is reinfected with malware within five minutes—repeated reinfections signal the presence of a rootkit or persistent compromise.

> 8. A user account trying to login to multiple resources within a few minutes from/to different regions—a sign that the user's credentials have been stolen or that a user is up to mischief.

> Time-sensitive item.

Detect, deny, and disrupt attacks

If we know what to look for, how do we detect and disrupt based on these indicators? We found that while countermeasures may surface the relevant data, three factors hold back response efforts:

• The full potential of existing countermeasures have not been activated (systems are kept at default settings or do not implement threat intelligence services).

• Relevant data has not been captured, retained, or shared.

• Companies with outdated SIEM, firewall, and endpoint protection lack real-time correlation and fine-grained rules—and this prevents them from elevating the criticality of key indicators.

Default deployment and laggard adoption

While products and services may go live with "secure by default" settings, no default setting reflects an optimal configuration for every business. In addition, the most advanced features may require upgrades that companies choose not to implement because of concern about changes affecting availability, status quo compliance, fear of false positives, or lack of testing resources. However, a "laggard" approach to feature adoption substantially cripples protection when the threatscape is changing so quickly. In business terms, the cost avoidance enabled by laggard feature adoption becomes cold comfort when an attack succeeds and carries millions of dollars in penalties, brand damage, and costs.

When it comes to attacks, companies need to lead, not lag. Companies using a real-time SIEM could detect threats in minutes, companies without one took far longer.
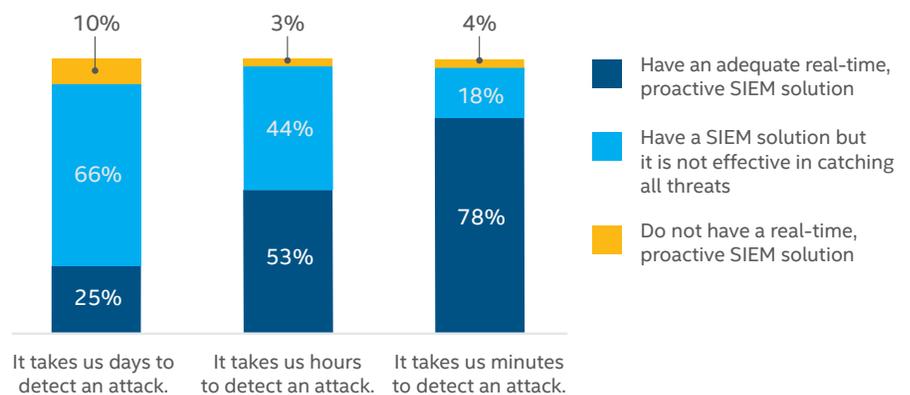

Figure 6. Evalueserve Question: While many factors affect detection speed, the adoption of real-time SIEM correlates with fastest detection.

Threat intelligence is another prime example of an area where adoption may lag, causing security effectiveness to suffer. Threat intelligence allows enterprises to learn from the experiences of others and block activity that might otherwise be difficult to identify as malicious. The defenses can block in real time if they can access up-to-date reputations for "known bad destinations" and other dynamic threat attributes.

Success Story: Government of New Brunswick

Jamie Rees, director of information assurance and chief information security officer, explains how SIEM has enabled his team to spend more time on strategic projects and less tending to workstations.

"Now, with the SIEM automatically pulling and correlating data from devices throughout the GNB ecosystem, the team is able to access the information instantly from its central location, generate detailed reports that identify events of concern that should be elevated, and then work remotely with the involved departments to quickly implement a fix."

When McAfee surveyed its customer base, organizations quantified substantial protection improvements when they activated McAfee Global Threat Intelligence services. Top benefits came with the network intrusion prevention system (IPS) and endpoint protections, but every protection saw at least a 20% bump up in protection—as assessed by the administrators themselves. When asked about the specific benefits, 29% called out a reduced time to detection, with an average improvement of five days.[4] This is just the most basic threat intelligence. Results improve when linked with correlation rules, risk scores, and local threat intelligence.
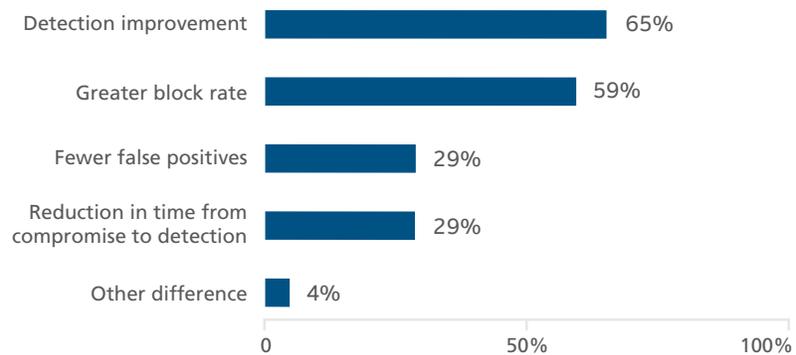
Detection improvement **65%**
Greater block rate **59%**
Fewer false positives **29%**
Reduction in time from compromise to detection **29%**
Other difference **4%**

0   50%   100%

Figure 7. Activation of McAfee Global threat intelligence has a material impact on detection of modern threats.

Data collection and aggregation

The second common hurdle is data collection, retention, and aggregation across functional boundaries—product silos and organizational groups. Although this situation can create a data morass if improperly managed, most organizations benefit from being selective. The key is consideration of what data is truly the most valuable for that organization. For this, security and IT teams need to work with one another and business partners to identify key assets, the baseline of appropriate use of those assets (such as applications, users, time of day, typical workloads), relevant indicators of attack for the asset's likely threats, and countermeasures in place.

For example, an accounting database server houses sensitive data, communicates using specific ports and protocols, has a finite set of approved applications and users who work within typical workdays, plus occasional peaks (end-of-month reporting, for example). As an SQL database running on Linux, attackers would attempt to exploit vulnerabilities in the database or the underlying operating system, which can be mitigated through countermeasures, such as application whitelisting, database activity monitoring, and network intrusion-based prevention. By documenting and baselining these characteristics, the IT team can set alerts for and act on any unusual behavior.

To bring this value home, business knowledge enables 50% of the indicators of attack by guiding understanding of "internal hosts." The second indicator of attack, for instance, requires visibility into internal hosts as well as correlation with unusual port/protocol combinations. If a database server suddenly starts communicating over port 80 via FTP (IoA 2), something is amiss. A baseline profile would reveal this misbehavior. If the database server starts to talk to a system in the DMZ (IoA 3), that could be the path out for valued data.

Success Story: Kroll

Kroll, a global risk management company, counts on automation to multiply the effectiveness of its security staff. Gene Cupstid, senior information security engineer at Kroll, explains:

"From my perspective, one of the most important SIEM capabilities is automated responses based on correlation rules. I love that I can have the SIEM proactively run a script or place a block on the IPS or even automatically kick off a virus scan. We can even have it lock down a host and take it off the network. All of those things can be automated given the right business intelligence and the right SIEM."

Correlation and rich rules

Finally, increasing attack sophistication justifies increasing response sophistication. As the survey showed with the dissatisfaction and slow response time of many SIEM users, legacy tools may: not have the required context, permit complex rules, or be designed to sustain this level of data and intelligence processing. Data aggregation, multiple-attribute logic, and conditional or multiple-step rules must be active for nuanced identification and action.

In practice, the correlation performed by a real-time SIEM could help factor in the consideration of frequency within a time window or activity outside a time window. Correlation rules can be used to detect off-hour malware detection, network scans by internal hosts communicating with multiple hosts in a short time frame, multiple alarm events from a single host or duplicate events across multiple machines in the same subnet over a 24-hour period, and repeated user login attempts from different locations. All of these indicators will be ignored by an individual countermeasure with binary (on/off) signature matching or enforcement actions.

Beyond mere monitoring, these clues enable proactive improvements to security posture. Administrators can refine risk sensitivity, set SIEM watch lists for recurrence, mine historical databases for similar and related events, and adjust rules and policies. And baselining of approved behavior helps too.

Simple Rules	Conditional Rules
Look for a signature for that file.	Alert if the same signature is applied more than *five* times within *five* minutes.
Block this type of communication.	Block this type of communication outside of working hours.

Appropriate automation

Our experience indicates that the fastest intervention—automated thwarting of a breach or disruption—builds on the aggregation and correlation of indicators of attack like these. High-speed logging, alerting, and contextualization by a real-time SIEM can assemble event data and elevate attack attributes earlier in an attack and help surface significant events to the security operations center (SOC). Of course, indicator of attack information can be collected and correlated manually, but that typically takes weeks and months and leaves your organization at the mercy of the attacker.

Once an attack is identified, workflows and thresholds allow selective use of automated responses based on risk tolerance, asset type, and threat. Automated responses might launch scripts and scans, update endpoint policies, or spark the IPS to quarantine compromised hosts. These automated workflows are not only more efficient, but they ensure consistent application of countermeasures and policy changes against specific threat categories.

Historically, false positives have made people leery of too much automation, but the volume and diversity of attacks require updated tactics. As a hybrid approach, individual automated correlations and their interpretations can be connected into workflows, with both manual and automated approval steps. This structure allows organizations to more speedily and consistently follow response policies and procedures, yet incorporate human and dynamic decision making.

Summary

An effective defense against advanced threats hinges not only on being able to detect pernicious intruders, but doing so in time to prevent significant damage to business operations and assets. This negative impact is the key variable in the risk equation: Risk = Threat x Vulnerability x Impact. By the time forensic analysts comb through mountains of security data looking for indicators of compromise (IoCs), their organizations may have already incurred losses.

This report aggregates front-line experience to present the key steps organizations can take to harden their infrastructure, improve their responsiveness, and actively disrupt targeted attacks by paying attention IoAs. It includes ways to learn from each interaction, enforce consistency, connect the smoke signals of an attack, and create an actionable, real-time picture of dynamic events.

A real-time SIEM is a significant enabler, since continuous monitoring and advanced analytics allow security managers to identify IoAs quickly and accurately. Integration can even catalyze instant action to contain and remediate the attack.

But the reality is that technology is not always the problem. Many of the countertactics mentioned here can be implemented with existing countermeasures and an integrated incident response program.[5]

The information most helpful to success can be recognized and mitigated today with adequate people and process and with technology organizations many have already deployed. The call to action for risk and threat managers is to focus on time management: improving their ability to detect, respond to, and learn from events as they unfold—thinking and acting within a timeline expressed in minutes.

This mature approach to proactive incident management as part of an overall risk management strategy offers the most agile and effective protection against targeted attacks.

About Intel Security

McAfee is now part of Intel Security. With its Security Connected strategy, innovative approach to hardware-enhanced security, and unique Global Threat Intelligence, Intel Security is intensely focused on developing proactive, proven security solutions and services that protect systems, networks, and mobile devices for business and personal use around the world. Intel Security combines the experience and expertise of McAfee with the innovation and proven performance of Intel to make security an essential ingredient in every architecture and on every computing platform. Intel Security's mission is to give everyone the confidence to live and work safely and securely in the digital world.

www.intelsecurity.com

1   http://home.mcafee.com/advicecenter/?id=ad_phishing_optpopnp

2   http://www.mcafee.com/us/resources/white-papers/wp-sans-incident-response-fight-back.pdf

3   Phone interview, July 30, 2014

4   McAfee Customer survey of Global Threat Intelligence Adoption, August 2014

5   http://www.darkreading.com/partner-perspectives/intel/drag-your-adolescent-incident-response-program-into-adulthood/a/d-id/1317162?