Modern IP theft and the insider threat

Mark Warren

Mark Warren, Perforce

With high-profile data breaches increasingly making the headlines right around the world, cyber-security has become a priority for company boards across virtually all business sectors. The whole topic of security has crept up the corporate agenda for various reasons, including concerns around intellectual property (IP) theft by employees or outsiders impersonating employees.

This is highlighted by recent PwC research into global cyber-security, which reported on the increasing incident of IP theft – data and other assets – and pointed out that most security incidents are caused by company insiders.[1] With the rapid transition of almost every industry to being driven by software, it's not surprising that the impact of cybertheft has more potential than ever for serious damage. The research also noted that investment is shifting from prevention, which is of course still important, to detection and reaction.

Far-reaching consequences

Clearly, leakage of IP has far-reaching consequences, damaging competitiveness, innovation and potentially leading to massive commercial losses. A US Department of Commerce report found that IP theft (all kinds, not just cybercrime) costs US companies $200-250bn annually, while the IP Commission Report puts this figure in excess of $300bn. The Organisation for Economic Development (OECD) estimated that counterfeiting and piracy costs companies as much as $638bn per year. On a wider scale, IP theft can damage economies and endanger state security.

IP breaches can be catastrophic for employees too. Financial losses from cyber-theft could cause as many as 150,000 Europeans to lose their jobs, according to a cybercrime report by Intel Security/McAfee in 2014.[2] In addition, the same report cited the example of a firm with 800 employees that, "had to cut its workforce in half after hackers stole its IP and a competing product appeared on the market."

Software development

Imagine being the network administrator at a computer games firm and finding out, when it's too late, that its code for its new launch was stolen by a rogue employee who's about to leave the company. Or being an automotive firm that's just found out that its top secret code associated with a new breakthrough in vehicle engineering is being mass-produced by a firm on the other side of the world, simply because someone managed to mimic a valid employee's ID.

"Enterprises are in a perpetual race against time to deliver the best products quickly and at higher quality than their competitors. Each product cycle generates vast amounts of mission-critical IP"

Given the scale of recent cyber-security breaches, there's a strong argument that the perimeter-only security model is not sufficient. Signature-based tools are simply unable to keep up with constantly changing new attacks. In addition, with Bring Your Own Device (BYOD) the

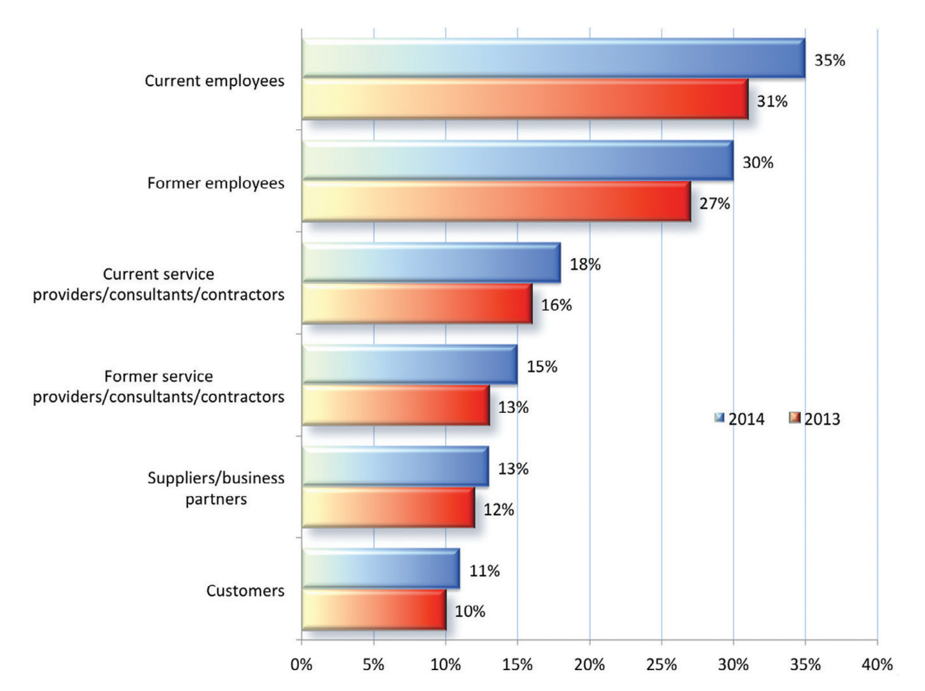

Figure 1: The top offenders of insider crime, 2013-2014. Source: PwC.

norm, data now routinely moves beyond the perimeter. Enterprises are in a perpetual race against time to deliver the best products quickly and at higher quality than their competitors. Each product cycle generates vast amounts of mission-critical IP, which can include initial product requirements, detailed engineering (hardware or software) specifications, industrial designs, source code, media, early prototypes, finished product, as well as business documents to help market and sell that product.

In such a software-driven era, security systems have to be able to address a wide variety of digital assets, created using many different tools by contributors with varying levels of technical ability and personal workflows. These assets may be stored in many different siloed repositories, not only making them harder to protect, but also impacting product quality and time to market. Preventing and tracking this kind of IP theft has traditionally been difficult and not something completely addressed by conventional security tools. This is why more organisations are turning to techniques such as behavioural analytics in the fight against IP theft, detecting and surfacing anomalies, such as unusual activities and applying algorithms that sort through all the noise. Before we take a closer look at those techniques, let's look at who perpetrates IP theft and the different types of attack involved.

Understanding the perpetrators

Hacktivists: Although the two most famous hacktivists, Bradley Manning and Edward Snowden, have garnered a major share of the headlines, these types of attacks are rare. They can be perpetrated by insiders, but can also originate from outside activist groups looking for social justice.

With social justice outlets such as WikiLeaks augmenting traditional news media outlets, it has become very easy for an insider to capture sensitive data and publish it while maintaining anonymity. If a company's corporate policies, products or projects touch on sensitive areas related to social, political or environmental issues, it must consider and defend against this attack motivation.

Criminal organisations: Organised criminal groups frequently use the Internet to commit fraudulent actions in the banking and financial system and the e-commerce sector. Primarily located in Eastern Europe and Russia, where law enforcement is difficult, these organisations have an underground marketplace where cyber-criminals can buy and sell stolen information and identities.

> ***"Studies consistently find that almost 60% of former employees have taken sensitive company data when they depart an organisation regardless of the reason why they left"***

The challenge security teams face is that these attackers will go after any data they can monetise, and IP that can be sold to competitors in foreign countries is becoming a favourite target.

Careless and compromised employees: Employees who move data to insecure locations in order to ease their work processes create risk by unwittingly exposing this data to external hackers or bad actors who work within a company, at supply chain partner companies or among contractors. A Cisco study entitled 'Data Leakage Worldwide: the high cost of insider threats' found that 44% of employees share work devices with others, 46% of remote workers admitted to transferring work files to home computers and 18% admitted to sharing passwords.

The compromised employee is the problem to which no one wants to admit, but it is still a very common form of data loss. The attacks are often long-term, moving small amounts of data over a long time.

Leaving employees: Employees leaving and taking sensitive data with them is a widespread problem. Studies consistently find that almost 60% of former employees have taken sensitive company data when they depart an organisation regardless of the reason why they left. One Symantec study found that 56% of workers believe it is okay to take data with them and use it at a competitor.[3] This includes not only customer contact lists, but also the IP and trade secrets related to the programs with which these employees were involved.

Also be very aware of disgruntled employees who are leaving; in September 2014, the US Department of Homeland Security put out an alert that it was seeing a significant increase in attacks of this category, which can cost businesses up to $3m per incident.[4]

State-sponsored cyber-espionage: China's People's Liberation Army (PLA) has developed a combat strategy called 'Integrated Network Electronic Warfare' that guides computer network operations and cyber-warfare tools with the goal of seizing control of an opponent's information flow and establishing information dominance. Analysts have long linked a unit in the Chinese military's 3rd Department to extensive cyber-espionage. This Chinese military unit has orchestrated attacks against 141 companies spanning 20 major industries.

Understanding the attack models

The impulsive attack: This is the preferred attack method used by employees leaving an organisation with sensitive data, as well as by insider hacktivists. (Snowden and Manning fall more into the long-term attack category.) This attack occurs in hours or in just a few days and is accompanied by significant anomalous activity. In many 'quick-action' incidents, attackers will access data they have rarely or never accessed, execute events that change or obfuscate that data and finally move large amounts of removable data to storage devices, personal machines or public cloud storage.

The below-the-radar slow attack: Much more common in government or corporate espionage, these attacks have been known to last for years. In one attack, a senior engineer from a major defence contractor stole sensitive US Navy data and went undetected for over three years. To remain undetected, these attacks often move only small amounts of data over a long period of time and most commonly out of the network through removable media or personal devices. These attacks are not just limited to insiders – it is common for outside attacks to also follow this model to avoid detection.

The outside or targeted attack: The outside attack, often called targeted attack or APT (advanced persistent threat), is a form of cyber-attack that is characterised by a high degree of technological and process sophistication mixed with a prolonged duration. Such attacks require a significant amount of resources to mount and are therefore usually sponsored by, if not directly controlled by, government or military organisations. They target specific, predefined organisations and particular data within them. Less common are insider or simple malware-based outside attacks. These attacks are very difficult to defend against with existing security tools because of their 'continuous, changing attack' model with the ability to find and bypass defences.

The twofold or inside/outside attack: Considering the level of sophistication employed by nation state-sponsored insider attacks and the growing number of targeted outside attacks by these same nation states, it's logical to conclude that the next big threat combines the strengths of the two.

> ***"For most large companies, it is not uncommon to deal with a staggering 100,000 alerts a day. This amount is overwhelming and creates a mindset known as 'alert fatigue'"***

Where targeted outside attacks take months to penetrate an organisation's defences, compromise one or multiple machines and manoeuvre to the targeted data, this same attack takes minutes for a competent insider. Yet when the insider is captured, it is a public embarrassment to the perpetrating nation state. The new scenario is simple: have the insider introduce the malware behind the defences and then relinquish control to an anonymous command and control (C&C) mechanism. With the malware present and obfuscated, the C&C server quietly and continuously extracts data out of the organisation.

Intelligence is the answer to security noise

With all of these perpetrators and attack models, there is little doubt that security teams face a huge challenge in protecting their company's IP. Most security teams have 10 or more different security tools deployed in hopes of detecting and preventing attacks, but that also means they have to manage those systems and the alerts they generate. For most large companies, it is not uncommon to deal with a staggering 100,000 alerts a day. This amount is overwhelming and creates a mindset known as 'alert fatigue'. In a recent *Wall Street Journal* blog, Gartner analyst Avivah Litan was quoted describing a client who receives over 135,000 security alerts a day.[5] As Litan aptly stated: "It becomes like the car alarms going off in a parking lot – no one takes them seriously because generally there are too many false car alarms."

The answer to the 'security noise' problem is intelligence – intelligence that can understand the context of an attack and accurately rate it in terms of risk priority. Dealing with 100,000 alerts is manageable when they are ranked in order of how risky they are, so the security team can clearly understand which threats needs to be addressed and in which order.

Threat detection and behavioural analytics

The biggest challenge with these attacks is that existing security technologies cannot detect them. And if attacks are somehow detected, this typically happens months after significant data breaches have already occurred. To successfully protect critical IP in this new threat environment, organisations must have a deeper understanding of what's happening with their important data, so that they can see and understand when it's truly at risk. The key to detecting this risky anomalous IP access behaviour is to identify the users, machines and projects involved in risky abnormal actions and then stop IP theft before it happens.

> ***"Using a weighted anomaly approach in combination with machine learning effectively minimises and, over time, reduces the noise and false positives that plague security teams"***

Statistically, human decision-making processes can be observed, measured and even predicted if tracked according to each person's unique decision-making patterns and risk-tolerance levels. Modern threat detection uses sophisticated mathematical models to do exactly that. It identifies that a pattern of behaviour has deviated from its norm, but also quantitatively measures the probability that the observed behaviour is risky. For instance, someone accessing a single important source code project more often than they have historically accessed it is interesting, but not as interesting (or potentially as risky) as someone accessing 10 important source code files that they have never accessed before. Such examples are weighted even higher if they are identified in close proximity with other anomalies involving the same entities (users, source code projects, source code files), including time of activity, volume of activity and movement of data.

New threat detection behavioural analytics find anomalies by:

- Comparing access patterns, data usage patterns and data movement patterns against historic behaviour.

- Determining similar user patterns across the environment and comparing behavioural patterns between users and groups of users.
- Detecting dissimilar patterns among members of the same project group or job role.
- Comparing individuals against the entire user group.

Weighted anomalies

Not only are these anomalies leading indicators of threat activity, they also require no foreknowledge or configuration to detect. Using a weighted anomaly approach in combination with machine learning effectively minimises and, over time, reduces the noise and false positives that plague security teams. By analysing the four types of attacks discussed earlier – impulsive, under the radar, outside and twofold – we find that behavioural analytics offer valuable insight:

- Impulsive: The comparison of users against their historic activities will easily surface the attack.
- Under the radar: Although a user may intentionally attempt to prevent detection, analytics are used to compare the user to his or her peers over time to detect and reveal indicators of attack. Under the radar is a consistent pattern that is very different from normal peer activity.
- Outside attack and twofold attack: In these cases where machines are compromised with stealth malware that is attempting to siphon sensitive information externally, behavioural analytics can interpret normal access and usage patterns for all users. Compromised machines or identities do not know what is historically normal and therefore deviate and create anomalies.

This data-centric internal threat detection approach uses behavioural analytics and their underlying math models to aggregate, correlate and measure the risk of each threat and then present a clear and accurate depiction of which threats need to be addressed in order of priority (see Figure 2).

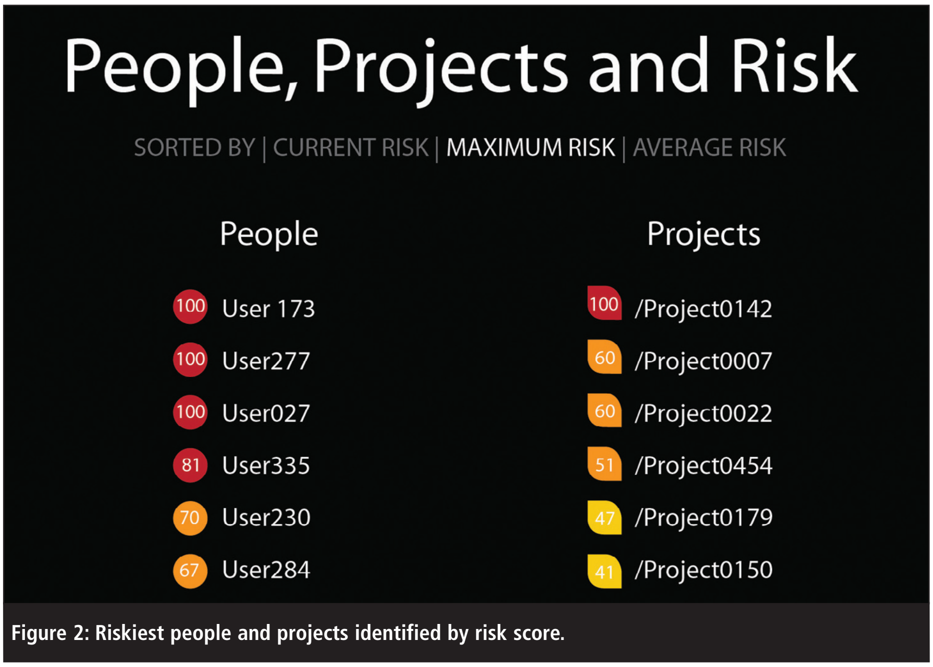

People, Projects and Risk

SORTED BY | CURRENT RISK | MAXIMUM RISK | AVERAGE RISK

People		Projects	
100	User 173	100	/Project0142
100	User277	60	/Project0007
100	User027	60	/Project0022
81	User335	51	/Project0454
70	User230	47	/Project0179
67	User284	41	/Project0150

Figure 2: Riskiest people and projects identified by risk score.

How threat detection works

Behavioural analytics is not new, but applying these proven methodologies for identifying and mitigating risk within security is a paradigm shift. Using the latest threat detection techniques can optimise the entire threat detection process, from data collection to analysis and reporting.

Advanced behavioural analytics are applied in two specific areas: behaviour risk and entity risk. Both sets of mathematical models influence each other, such that behaviour risk models incorporate the entities involved, and entity risk scores are influenced by risky behaviours that they are involved in.

Behaviour risk scores are applied to all observed behaviours, across different behavioural vectors. For example, when a user takes a certain amount of data from a project, that volume is compared to that user's historical baseline, as well as other similar users' baselines. This provides multiple comparison points, which refines the overall behaviour risk score.

Figure 3 is a simple three-event example that shows the relationship between behavioural and entity risk models and how entity risk scores change over time. As 'John Sneakypants' executes three events, the anomalous nature and riskiness of each event create higher behavioural risk scores. Entity risk scores are low at first, showing minimal danger for the user and the project that the user has accessed. As each event occurs, behaviour risk scores climb, and associated enti-

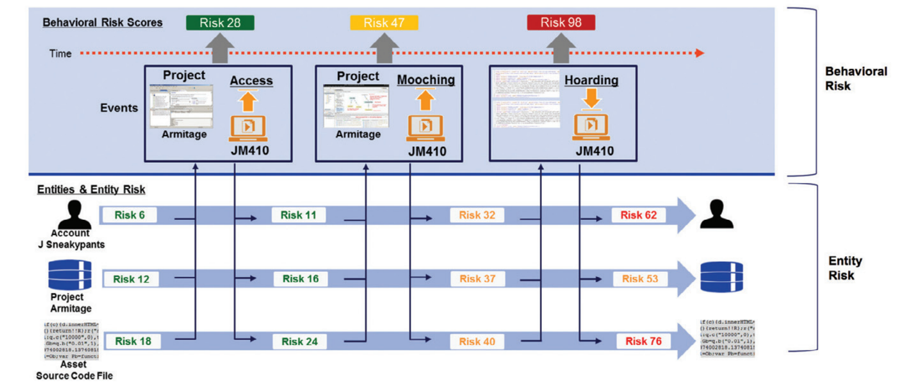

Figure 3: The relationship between events, behavioural risk and entity risk.

ties' risk scores also climb. The analytics engine captures and surfaces the threat across the events as well as the entities.

Powerful method

Behavioural and entity risk models represent a powerful method to remove the noise that currently overwhelms security teams.

Suppose John Sneakypants was detected accessing an important software project that he does not normally access based on his historical access patterns. In addition, his peers do not normally access this important project. This action may be suspicious, but it could also be a false positive if John has had a recent role change or has been assigned to a new project. But suppose John also accessed this file at a time of day that he was never active before and that he also just took files from another source code project that had been inactive for months. Plus, the overall volume of data he's taking from multiple projects is highly unusual. The more risky events occur involving the same entity within a short time frame, the less this feels like a false positive, leading to increased urgency to investigate John (see Figure 4).

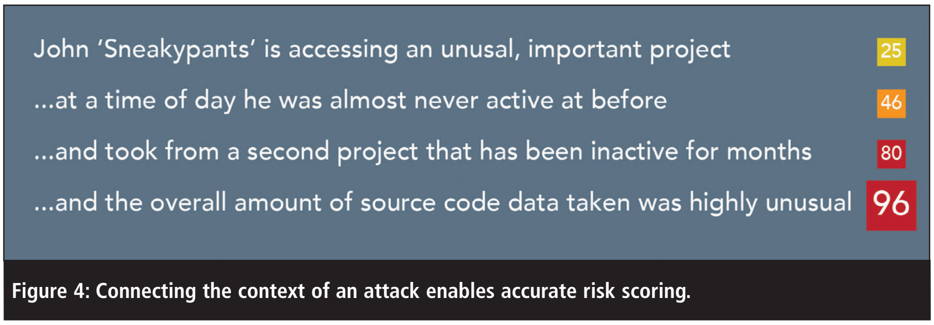

John 'Sneakypants' is accessing an unusal, important project 25

...at a time of day he was almost never active at before 46

...and took from a second project that has been inactive for months 80

...and the overall amount of source code data taken was highly unusual 96

Figure 4: Connecting the context of an attack enables accurate risk scoring.

Behavioural analytics illuminates not single events, but patterns and relationships created by the behaviours of users. This enables the focus on actual threats, while tuning out the uninteresting noise that overwhelms security teams. This new approach can vastly improve an organisation's ability to quickly determine the root cause of a threat and respond proactively before critical data is compromised.

> ***"Actionable reports can lower the cost to investigate an incident and accelerate legal action and/or adjustments to policies or security tools to prevent or reduce the risk of a future breach"***

Capturing the events and relationships is also the key to lowering the cost and time of forensic investigations. By capturing the relationships between people, behaviours and projects, an investigation can quickly and accurately identify the information that defines the risk or threat down to the user, project and time the behaviour occurred. Because all behaviour is tracked and captured, the events that led to the threat or incident are immediately available. The interface enables these relevant activities to be pinpointed, compressing the time it takes to determine what happened, who was involved and which projects are affected. Actionable reports can lower the cost to investigate an incident and accelerate legal action and/or adjustments to policies or security tools to prevent or reduce the risk of a future breach.

The experience of one unnamed $20bn manufacturer illustrates the process perfectly. It had become a victim of insider IP theft and, having already spent $1m with a traditional security vendor over 12 months, took a modern data-led approach to identify all the perpetrators. By using log data from 20,000 global developers from 30 days of activity – totalling over 9 billion events – and applying machine learning and analytics the internal threats were uncovered in just days. Not only did it identify the two known rogue engineers, but it also identified 11 other previously unknown thieves that had been stealing large volumes of information.

Conclusion

Today's cyber breach headlines highlight the sophistication of attacks and the failure of major organisations to protect their sensitive IP. Traditional perimeter-based security, as well as the current generation of security tools, has not succeeded in mitigating such attacks. Security teams are simply too overwhelmed.

To effectively detect and mitigate advanced threats focused on stealing critical IP such as source code, companies must look to new technologies that support the creation of data-centric internal threat detection. This strategy applies intelligence closer to the sensitive data it is intended to protect and can prevent advanced attacks or insiders from stealing this data even when perimeter and network defences have been defeated. Data-centric threat detection combines continuous monitoring of an IP repository with advanced behavioural analytics, dramatically decreasing the risk of IP theft.

This approach reduces the overall cost and complexity of a threat detection and data protection programme, while increasing a security team's ability to reduce risk and surface actual threats to the entire organisation. Enterprises can now implement a principal internal defence strategy with technology specifically designed to monitor any corporate IP repository, surface risks and threats and alert security teams with actionable information to stop an attack before data is compromised.

About the author

Mark Warren is European marketing director of Perforce Software, with additional responsibility for solutions marketing. Worldwide, the version management and code collaboration portfolio from Perforce Software is used by thousands of customers, including Salesforce.com, Nvidia, Samsung, and EA Games. Warren has over two decades' experience in the soft-

ware industry with roles as a provider and consumer of advanced development tools.

References

1. 'The Global State of Information Security 2015'. PwC. Accessed May 2015. www.pwc.com/gsiss2015.
2. 'Net losses: estimating the global cost of cybercrime'. McAfee/Centre for Strategic and International Studies, Jun 104. Accessed May 2015. www.mcafee.com/us/resources/reports/rp-economic-impact-cybercrime2.pdf.
3. 'Symantec study shows employees steal corporate data and don't believe it's wrong'. Symantec, 6 Feb 2013. Accessed May 2015. www.symantec.com/about/news/release/article.jsp?prid=20130206_01.
4. 'Increase in Insider Threat Cases Highlight Significant Risks to Business Networks and Proprietary Information'. US Department of Homeland Security, Public Service Announcement, 23 Sep 2014. Accessed May 2015. www.ic3.gov/media/2014/140923.aspx.
5. Hickins, Michael. 'Target CIO takes the fall'. Wall St Journal, CIO Journal, 5 Mar 2014. Accessed May 2015. http://blogs.wsj.com/cio/2014/03/05/target-cio-takes-the-fall/.

Encapsulating mobile security

Ofir Agasi, Check Point

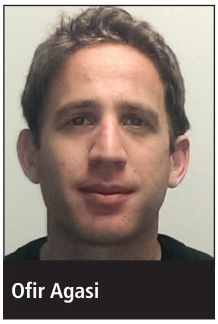

Ofir Agasi

Are there many businesses that are *not* using mobile devices? Recent research by Enterprise Strategy Group (ESG) showed that 87% of enterprise organisations say mobile computing is either 'critical' or 'very important' in supporting their business processes and employee productivity.[1] But the benefits of mobility often come at the cost of security.

Check Point's third annual mobile security survey found that, in a majority of firms, mobility is racing ahead of effective device and data protection.[2] Of over 700 IT professionals surveyed about mobile usage in their companies, 72% said the number of personal devices connecting to their networks had more than doubled in the past two years; 82% expected mobile security incidents to grow over the next 12 months, with higher remediation costs; and 42% noted that such incidents had already cost their organisations more than $250,000.

The reasons why mobile security is proving so challenging were highlighted by the ESG research report. The issues identified were:

- **Protecting data at rest and in motion:** 43% of security professionals said it is a problem to protect confidential data accessed from a mobile device, and 41% say it is challenging to protect sensitive data stored on a mobile device. Mobility creates blind spots in which the security team can't easily monitor sensitive data.
- **Enforcing security policies:** Most security policies were originally created with PCs and wired Ethernet ports in mind. Often, enforcing policies and delivering protection on mobile devices is approached by implementing new tools and infrastructure – adding extra management complexities for IT teams.
- **Integrating mobile security into existing security processes:** As organisations create a mobile security overlay infrastructure, it becomes increasingly difficult to maintain consistent policies, coordinate enforcement actions, or to monitor users and devices across the network.

Disparate solutions

It is this mix of challenges that makes mobile security so difficult to address. Various disparate solutions have attempted to address mobility and security, but none provides a complete solution. Enterprise mobile management (EMM) solutions manage device configurations, but do not secure business data and documents in uncontrolled environments. Similarly, point mobile products that focus on specific areas of security do not integrate with the organisation's corporate policies or IT infrastructure.

> ***"Mobile malware is becoming an enterprise threat vector, as attackers look to target ever-growing mobile estates"***

What's required is a modular, integrated approach that delivers functionality to address the three main mobile security problems. The required functionality is:

- **Threat management:** mobile malware is becoming an enterprise threat vector, as attackers look to target ever-growing mobile estates. Organisations should prepare for this by deploying controls and monitoring capabilities for mobile threat prevention, detection and remediation.