bjss

Moving to Analysis-Led Cyber-Security

A BJSS WHITE PAPER

Why do sophisticated cyber attacks succeed against today's security systems? Or why is it that they can be difficult, or even impossible, to detect, let alone prevent?

This BJSS White Paper aims to provide a summary of trends in sophisticated cyber-attacks, explain which industries are most affected and introduce risk management frameworks from mature security disciplines including physical security, fraud detection, and counter-terrorism intelligence.

The standard defensive toolkit relied upon by the enterprise security professional – anti-virus, pattern-based intrusion detection/prevention systems, firewalls, patching – remains vitally important. However, sophisticated attacks can penetrate all of these defences and are increasingly affecting a growing number of organisations worldwide. These trends are likely to continue. In order to cost-effectively mitigate risk, security teams within enterprises should learn lessons from some of the most mature security fields – physical protection, fraud prevention and counter-terrorism intelligence – and extend their focus from prevention to intelligent proactive and reactive risk management.

Don't Panic

It is important not to be alarmist. The vast majority of attacks an organisation is likely to face are still opportunistic, automated 'script kiddie' attacks. These can largely be mitigated by (seemingly) simple preventative measures such as quickly applying security patches, following OWASP guidelines for custom-built applications, or by deploying basic measures such as firewalls, intrusion prevention systems, drive encryption and anti-virus software.

In practice, even ensuring that these basics are implemented consistently and proven to be working can be very difficult for many security groups – especially those without strong automated IT configuration management systems across their IT estate. It is important that CIOs, CTOs and CSOs retain their focus on this area.

However, the Snowden revelations clearly illustrate that some of the most determined attackers – well-funded nation-state intelligence services – developed capabilities several years ago that far outstrip most currently deployed defences. They are interesting for information security professionals because they provide a window into high-end offensive cyber operations, a sphere of activity that is normally closed to the outside world.

Although the information that was leaked related to the 'five eyes' Western intelligence organisations (UK, Canada, Australia, New Zealand, and the US), and it is impossible to be sure whether all of the released information is accurate and is not taken out of context, it would be naïve to assume that many states and state-linked or sponsored groups do not have similar capabilities. Also, the fallout from reports that interception of inter-datacentre links of major Internet players has taken place has shown that even US companies are not immune and can become victims of 'friendly fire'. This has resulted in a huge impact to public reputation and trust, leading to consumers questioning the motives and loyalties of their favourite technology and online brands.

Target-Rich Environment

The revelations have also reinforced that a wide variety of organisations may be active targets for nation state level attacks. These include the 'usual suspects': communications service providers (whether telcos, ISPs, Internet backbone providers or social networks), government departments and defence companies.

In addition to these, it is likely that there are preparatory and ongoing compromises of critical national infrastructure providers, including big players in the financial services, transport, energy/utilities and healthcare industries. These are often invisible because they have not yet been – and may never be, unless a true cyber war occurs – actively exploited.

Perhaps less obviously, key suppliers serving all of the above industries are also at risk, as disruption of the supply chain is standard military doctrine.

The media industry is also a target: it can help to identify dissidents. Strategic compromise of websites that are industry 'watering holes' is becoming an increasingly common channel for spreading malware within specific industries. Third parties often have a more relaxed security posture and therefore represent a weak link or an easy point of entry for a targeted attack. Crowdstrike's view of the affected industries in 2014 is illustrated in Figure 1 below.

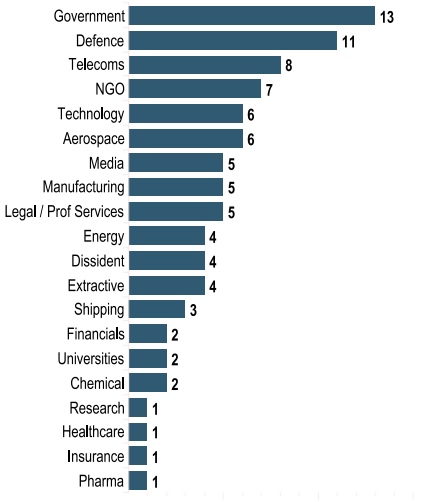

Figure 1:
CrowdStrike breakdown of industry attacks based on observed adversary activity in 2014.

Organisations which feel that nation threats can be ignored, because the chance of exploitation is so remote and many companies will be in the same boat, should note that offensive military cyber cells are not perfect. They themselves may be infiltrated by cyber-criminals or rogue members of staff intending to make money 'on the side' by either selling IPR or engaging in other forms of industrial espionage.

An Industry-Wide Challenge

While state actors are by far the most prolific originators of sophisticated attacks, others are also capable of pulling them off. A sophisticated attack is distinguished from more mundane ones by some or all of the following factors: a high degree of technical expertise; a willingness to devote substantial time and effort to each target; a methodical approach; and a willingness to use a blended approach of remote exploitation, social engineering, physical attack methods and insider compromise to achieve the objective. Microsoft defines the attacker and attack types as 'determined' and 'targeted' respectively.

Even those who aren't targeted by a nation state can find themselves falling victim to 'medium-level' hacks. These take the form of a diverse range of threats such as organised crime groups seeking personal information of customers, corporate espionage firms seeking intellectual property (IP), activists, journalists, disgruntled insiders or even skilled individuals who 'hack for profit'. While current security monitoring and response solutions will catch the majority of these attacks in time to prevent serious problems – it is only effective where the capacity exists within the security team to follow up alerts.

This presents a challenge for many organisations. It is difficult for a small team of security staff to stay on top of an increasing flow of alerts that come from even a small number of monitoring solutions. In practice, many security-conscious enterprises face a proliferation of alerts. If security teams attempt to triage and only look at the most important events, they will leave themselves vulnerable to attacks that are designed to 'slip beneath the radar'. More advanced hackers will research (often with the aid of insider assistance or social engineering) the characteristics of a target, and then adapt their behaviour to camouflage themselves among the noise. At times, and as a simple diversionary tactic, they will unleash a flurry of crude and ineffective but noisy attacks. The approach is simple: a deluge of information from monitoring systems blinds instead of illuminates.

Even if today's organised crime attacks can be caught using traditional techniques, the leaking of capabilities designed to defeat them from the advanced intelligence community provides an opportunity for organised crime to leverage these methods for their own purposes. In the words of consultant and cryptographer Bruce Schneier, "*today's NSA secret techniques are tomorrow's PhD theses and the following day's cybercrime attack tools.*" As a result, all security professionals should be aware of these techniques and the ways in which associated risks can be mitigated.

Risk Management in a Dangerous World

Stopping advanced attacks is hard. In fact, many experts such as Microsoft believe that targeted, well-resourced attacks will more often than not be successful – even for those with a well-developed security posture. Defensive security is only as strong as its weakest link – while attackers only have to be lucky once, defenders have to be lucky all of the time. Deploying truly advanced cyber defences is difficult, expensive, and can often impact on business agility. It is not feasible, desirable or economically sensible for any organisation to try to defend against all possible attacks.

The comfortable illusion that enterprise IT systems can be 'secure' is rapidly beginning to evaporate. Week after week, some of the most technically able organisations in the world are admitting to breaches. Next-generation preventative security will need to be less 'one-size-fits all' and more tiered and tailored. Data will increasingly need to be classified according to its sensitivity, and different controls and techniques applied that are commensurate to the risk.

This implies that a level of risk will have to be accepted in exchange for competitiveness and productivity, which is not a trade-off that many organisations are experienced in making in the IT security area. However, this does not mean that risk cannot be effectively managed. Historically, information security organisations have focused largely on prevention – but this is only a part of the risk management lifecycle. If attacks cannot be prevented cost-effectively, they can often be deterred or contained.

Maintaining a deterrent involves the prevention of simple attacks, credible detection and response capability for more sophisticated ones, and ideally an ability to carry out attribution of an attack. The aim is to force the attacker to increase the resources they must apply to carry out a successful attack – whether those resources are computing power, time, reputation, diplomatic goodwill or fragile tactical capabilities such as zero-day exploits. By raising the attacker's costs, some attacks simply become uneconomical. In other cases, their efforts might be redirected towards a cheaper target.

An additional benefit of effective detection and response is that it can also minimise the reputational and financial losses stemming from successful attacks, while also increasing the resilience of the defender.

An example of how attacks can be mitigated by investing across the risk management lifecycle is Anderson's chapter on physical protection in his classic textbook on security engineering. In section 11.2, Anderson describes how a physical security system is typically split up into a number of phases and elements. This is illustrated in Figure 2 below.

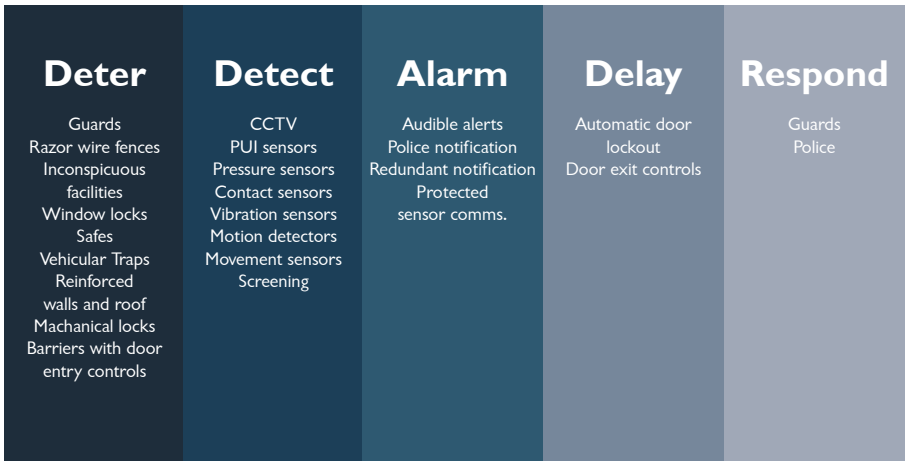

Figure 2: Physical Security Elements

Analysis-Led Cyber-Security

As the shift to detection and response gathers pace, it is important to realise that most teams will never have time to wade through every single alert generated by each individual system. One research study compiled by the Ponemon Institute found that 40% of financial services organisations, 52% of manufacturing organisations, and 59% of government organisations felt that they had "insufficient in-house personnel or expertise to analyse anomalous and potentially malicious traffic in [their] networks" – and that is with current capabilities, which will struggle to detect more subtle attacks.

Some of the problems faced by analysts in the relatively new field of cyber-defence bear great similarity to those faced by analysts in more mature domains – notably those working in areas such as financial fraud detection and counter/anti-terrorism intelligence.

In the financial fraud area, models of the fraud management lifecycle were described over a decade ago – one example is articulated by Wesley Kenneth Wilhelm in the *Journal of Economic Crime Management*. This sets out a framework for fraud management. It describes the key role that analysis plays in this lifecycle and the need to focus on all stages in a balanced way. The lifecycle has marked similarities with Anderson's physical security lifecycle.

Many of the current breed of SIEM tools focus on automated, predictive analysis of attacks using data mining techniques as a way to reduce the load on analysts. This idea is superficially attractive, and can be effective at spotting and preventing simple cases. However, it is worth cautioning that in fraud detection, similar tools have been found wanting against advanced adversaries such as organised criminals. As described by Wilhelm, these adversaries are adaptive, and will change behaviour in order to evade specific detection tactics. Automatically derived rules tend to perform poorly when compared to those used by human analysts because they tend to focus on more robust indicators of fraudulent behaviour.

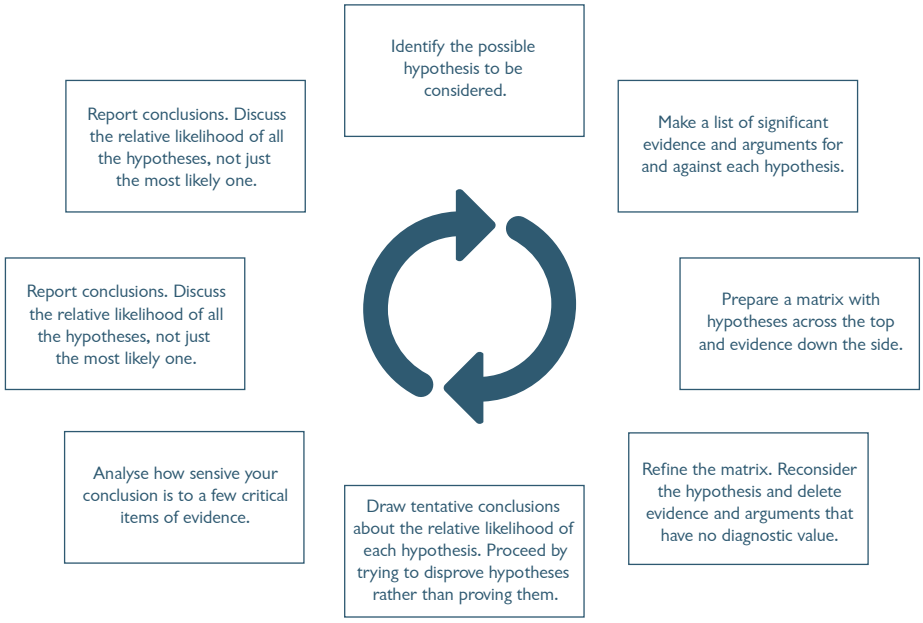

Figure 3: Analysis of competing hypotheses.

Counter-Terrorism

In counter-terrorism intelligence, data mining has been even less successful. A report by the Cato Institute illustrates why. There are relatively few terrorist incidents each year, so data that can be used for training is scarce. Any inaccuracy in detection rules means that false positives come to dominate available investigation time. The report states that the lessons learned post the 9/11 attacks are to not use data mining, but to "enable investigators to efficiently discover, access and aggregate relevant information related to actionable suspects". Sophisticated cyber-security attacks are similarly infrequent, so the lessons learned from counter-terrorism should be the same.

By contrast, intelligence analysis is a mature and proven field. Wikipedia defines it as: "the process of taking known information about situations and entities of strategic, operational, or tactical importance, characterizing the known, and, with appropriate statements of probability, the future actions in those situations and by those entities. The descriptions are drawn from what may only be available in the form of deliberately deceptive information; the analyst must correlate the similarities among deceptions and extract a common truth." Analysis of a determined cyber-attack has much in common because a well-executed attack will rarely leave a 'smoking gun'. Adversaries who have successfully managed to penetrate a network can often largely cover their tracks, or even leave 'false flag' information in an attempt to throw security analysts off the scent and mislead them as to what type of adversary they are dealing with.

Competing Hypotheses

One of the key tools of the intelligence analyst is analysis of competing hypotheses, as defined in Chapter 11 of Heuer's 'Psychology of Intelligence Analysis' and illustrated in Figure 3 above.

This formalises a structured method – which most good analysts will unconsciously carry out in any case – for analysing uncertain situations with potential deception present.

Another lesson from the intelligence world is that, due to the attacker's advantage of only having to find one weak point, being a passive defender is rarely adequate. To start to counter this, the military has a well-developed concept of 'kill chains'. These are "models of the systematic processes used by aggressors to target and engage an adversary to create desired effects". They can be applied in several areas: for example, Ashton Hayes describes a terrorist kill chain in order to provide a framework for mitigating against unknown terrorist attacks, emphasising a range of tactics across all stages combined with hardening of defences in a 'just in time' fashion.

Hutchings, in a paper published by Lockheed Martin, extends the kill chain concept to APT attacks. Here he provides a categorisation of tactics that can assist with detection, denial and disruption of cyber-attacks at each phase, but also those that can degrade attacks and deceive the attacker – as shown below in Figure 4. Again, analysis followed by responsive investment in defence is important. By their nature, APT attacks tend to play out over a long time, so this approach can be very effective.

Reidy presents an insider cyber-security threat kill chain derived from lessons learned at the US Federal Bureau of Investigation (FBI). This differentiates between pure APT attacks and insider threats – although it is of course possible that a very sophisticated attacker may use a combination of both. He reiterates that traditional security tools cannot help with an insider threat, and stresses the need for contextual data – such as HR and personnel data – to identify risks, in combination with more complex indicators. The FBI uses relatively limited network-level data, and takes care to baseline detection models for data exfiltration at the individual level.

Frustration and Deception

One way of deterring and detecting advanced attacks is to use frustration and deception techniques. Two examples are mentioned in the APT attack kill chain described earlier: tarpits (fictional sections of IP address space constructed by tools that accept but do not respond to network connections, or fake a variety of different

Phase	Examples	Detect	Deny	Disrupt	Degrade	Deceive
Reconnaissance	Research, identification and selection of targets, often represented as crawling Internet websites such as conference proceedings and mailing lists for email addresses, social relationships, or information on specific technologies.	Web Analytics	Firewall ACL			
Weaponisation	Coupling a remote access trojan with an exploit into a deliverable payload, typically by means of an automated tool (weaponiser). Increasingly, client application data files such as Adobe PDF or Microsoft Office documents serve as the weaponised deliverable.	NIDS	NIPS			
Delivery	Transmission of the weapon to the targeted environment. The three most prevalent delivery vectors for weaponised payloads by APT actors, as observed by the Lockheed Martin Computer Incident Response Team (LM-CIRT) for the years 2004-2010, are email attachments, websites and USB removable media.	Vigilant User	Proxy Filter	In-line AV	Queueing	
Exploration	After the weapon is delivered to victim host, exploitation triggers intruder's code. Most often, exploitation targets an application or Operating System vulnerability, but it could also more simply exploit users themselves or leverage an Operating System feature that auto-executes the file.	HIDS	Path	DEP		
Installation	Installation of a remote access trojan or backdoor on the victim system allows the adversary to maintain persistence inside the environment.	HIDS	"chroot" jail	AV		
C2	Typically, compromised hosts must beacon outbound to an Internet controller server to establish a C2 channel. APT malware especially requires manual interaction rather than conduct activity automatically. Once the C2 channel establishes, intruders have "hands on the keyboard" access inside the target environment.	NIDS	Firewall ACL	NIPS	Tarpit	DNS Redirect
Actions on Objectives	Only now, after progressing through the first six phases, can intruders take actions to achieve their original objectives. Typically, this objective is data exfiltration which involves collecting, encrypting and extracting information from the victim environment: violations of data integrity or availability are potential objectives as well. Alternatively, the intruders may only desire access to the internal victim box for use as a hop point to compromise additional systems and move laterally inside the network.	Audit Log			Quality of Service	Honeypot

Figure 4: APT attack kill chain and courses of action.

Operating Systems and services) and honeypots (fictional hosts, often with fake data residing on them that appear attractive to a hacker). These are described in detail in Strand and Asadoorian's book 'Offensive Countermeasures: The Art of Active Defence'. Tarpits use the fact that an attacker must probe for weaknesses in a network by slowing them down and setting off alarms. Honeypots provide the attacker with tempting but ultimately misleading information and, again, set off alarms.

The reason active defences are so powerful is threefold. Firstly, they reverse the attacker's advantage. In the same way that a single weakness in defence can result in a breach, a single mistake by the attacker will set off an alarm. Secondly, the false positive alarm rates should be very low, so response can be quick and decisive, without requiring extensive analysis. Thirdly, they interfere with what Strand and Asadoorian refer to as the 'Observe, Orient, Decide, Act' loop, slowing the attacker down, making them more cautious and potentially leading them in completely the wrong direction. This gives defensive security analysts more of a chance to respond to and frustrate the attacker, increasing the likelihood that they will make further mistakes.

Another type of active defence is that which helps with attribution. Some of the files served from honeypots – such as PDFs, Office documents, or Java applets – may contain 'phone home' functionality that is designed to reveal the attacker's true IP address, even if they are attacking via a proxy or anonymising service such as TOR. By attributing an attack to a particular area or group, it may be possible to better prioritise both incident response as well as an investment to shore up defences.

The experiences from the financial fraud and intelligence domains suggest that a successful approach to defending against advanced cyber-attacks is to provide technology that supports the entire risk management workflow.

Wherever possible it should automate tedious tasks, but also recognise that the investigative skills of expert analysts are key to managing complex security risks. Knowledge gained through analysis should be used to create an adaptive, responsive defensive system and to help recognise repeated attacks from the same attacker via their modus operandi and active attribution - analysis-led cyber-security.

Technology for Analysis-Led Cyber-Security

A natural question is how technology to support analysis-led cyber-security relates to and is supported by traditional security tools and techniques – in particular SIEM systems. There is certainly an overlap in capability, but the emphasis tends to shift from fully automated detection of security issues – which, as discussed previously, is rarely effective against a sophisticated attacker – to shortening the cycle time for human analysis of competing hypotheses.

Typically, the architecture and process will be data driven. Technology which supports analysts should be capable of integrating data from the obvious IT sources, including firewall (network and client-based), HTTP proxy logs, NIDS/ NIPS, HIDS, and AV alerts. However, it should also be able to ingest diverse contextual information which is not directly related to alerts. This should include actual detected configuration changes across the IT estate at network, server and client level, as well as planned change data (typically from a CMDB) that can be correlated automatically or by a human analyst to reduce false alarms. Some systems will also utilise specialist host agent logs recording process start/stop/new events and port usage.

Information on both successful and unsuccessful logins should be integrated from sources such as AAA infrastructure and application server logs. These are most useful when combined with data from systems such as physical access control (eg, door and site entry card systems), email 'out of office' notifications, calendar information, vacation authorisation systems and VPN logs. This allows analysis of whether there are accesses to server resources when the user does not appear to be present at an authorised endpoint, or when that endpoint is in an unusual location – for example, from abroad when the user does not appear to be on vacation or work travel, or from IPs at different ends of the country within the same three-hour period. The unreliability often typical of contextual sources means that an individual unusual login will likely be a false positive. However, a pattern of activity over a period of time is likely to indicate either a compromise or, at the very least, a member of staff persistently working around access control mechanisms, sharing passwords etc – which is typically interesting to a security team in its own right.

Most solutions will also be capable of efficiently capturing and retaining large amounts of metadata and payload data for information entering and leaving the enterprise and for internal traffic accessing key systems. Retention of this data for as long as possible is crucial because it may only become evident days or weeks after the event that a particular website has been compromised. At that point, it is important to be able to see which endpoints and users may have accessed it – and potentially, what they accessed. Even accesses to commonly permitted sites such as Facebook may be covert channels for botnet command and control. Integration of industry-wide external threat intelligence sources is important as a part of this.

Social Engineering

Attackers commonly use social engineering as a way of gaining access to accounts and protected resources. Helpdesk log information may provide evidence of this: for example requests to reset user passwords or requests to provide access to sensitive systems or shared folders. HR databases may provide valuable information to assess potential insider threats – from simple length of employment, through information gained during pre-employment screening that may indicate financial stress (eg, results of credit checks) through to active grievances.

Technology can automate some of the work that analysts would do to collate and link all of this information together and to search for patterns of concern, effectively generating hypotheses that analysts can explore in greater detail. Usually, this will involve linking events to entities such users, endpoints, websites, and IP addresses. Rules to detect risky events, whether generated internally or from existing SIEM-type solutions can then be 'rolled up' around entities to generate risk. High value intelligence such as accesses to honeypots can also be included as a high-scoring risk event.

The combination of multiple events helps to reduce the number of false positives. For example, several users may have visited the same website. Hostbased intrusion detection then picked up configuration changes on their endpoints. Regular polling of a particular Facebook page then takes place, which could be indicative of a command and control (C&C) channel, and finally users start accessing highly sensitive resources with unusual frequency. The combination of all of these indicators can be rolled up to the original website via graph analysis techniques to provide a high risk score to the website, even if no threat intelligence has yet been received. Another benefit is that alerting thresholds can be lowered on systems raising alerts, with aggregated, summarised information displayed to analysts. This lowers the risk of false negatives.

Because there is a limit to analyst time, tools should provide fine-grained prioritisation – usually via a continuous risk score, rather than just high/medium/low categorisation – providing a 'stack of work' to work down. It is important that analysts are presented with a clear summary of why particular entities or events have been flagged. This may, for first level responders, be via free text description of rules that have been triggered. However, first level responders may also decide that an attack is potentially worrying enough, warranting its forwarding to more specialist threat analysts.

Typically, these analysts are specifically trained in analysis techniques. They tend to rely heavily on visualisation and ad hoc searches to find additional contextual information, and tools that record their workings and findings in a format that is easy to share with other analysts. Investigation workflows tend to be ad hoc, so it should be possible for analysts to push actions and intermediate work products to other analysts or groups throughout the course of an investigation. Investigative case management software often provides a good basis for this.

Searching for Events

Event searching is critical to enable analysis to quickly find the information necessary to test hypotheses. Having to open multiple applications to run searches, or worse, manually request it from someone with access, severely impacts the analytical process. Ideally, searches should be capable of handling faceted drilldown on arbitrary categories, free text search, geospatial and value querying. However, there is a balance to be met between the cost of retaining data with sufficient indexing to allow this type of advanced query versus the benefit achieved. One way to resolve this is to retain older data with only basic indexing (for example, on time and key fields such as source/target IP address), but have the functionality to pull data back into richer indexes on demand.

Visualisation techniques often include link visualisation, timeline visualisation, geospatial visualisation and graphing. Figure 5 below shows typical examples of visualisations from Yarochkin's presentation 'Hunting in the Shadows: In Depth Analysis of Escalated APT Attacks'. Being able to auto-generate this type of visualisation from identified risk events or data located through searching can significantly aid the analytical process. It may be possible to turn some visualisations that can be automatically generated into live dashboards that can be displayed in security operations centres for ongoing situational awareness.

Finally, the tools should be capable of taking into account feedback from analysts based on the current operational situation – for example, emphasising or de-emphasising the importance of particular detection rules – and being rapidly configured to detect new patterns of behaviour corresponding to the shifting modus operandi of attackers. Most solutions will also have a capability to share key information on attacks (in a suitably anonymised form) to a central threat intelligence service to improve detection across multiple customers.

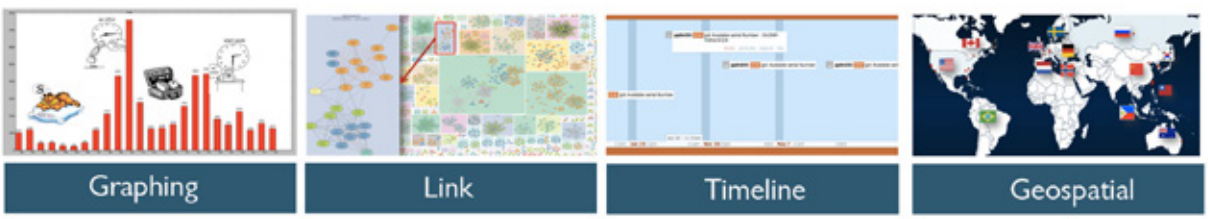

Figure 5: Visualisations from an analysis of an APT attack.

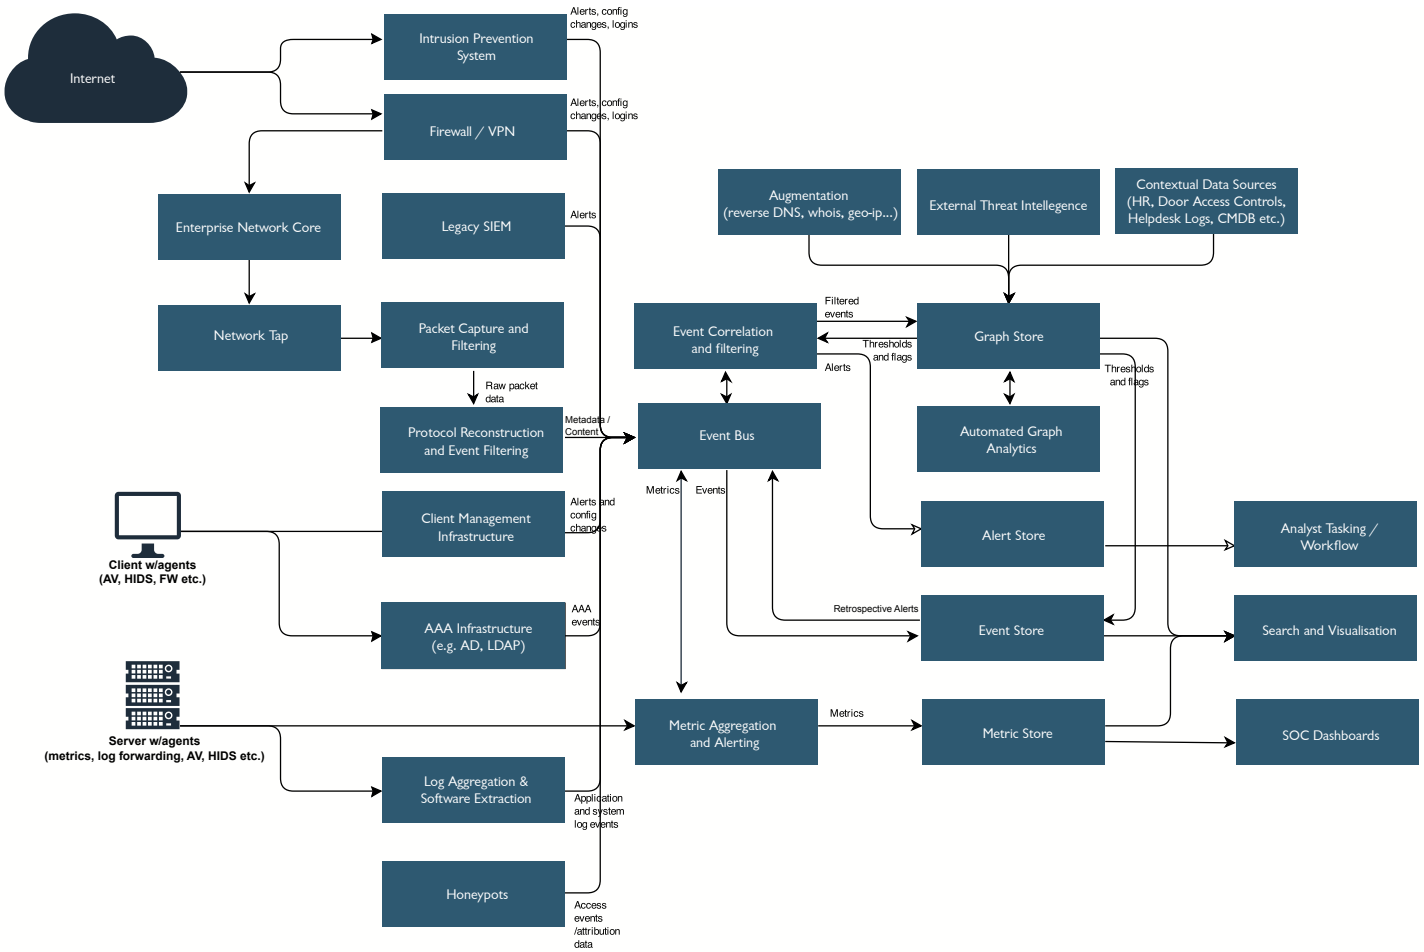

Figure 6: Reference architecture for analysis-led cyber-security

A Reference Conceptual Architecture

Given the wide scope of functionality that may be required for efficient analysis-led security, it is useful to be able to set technologies that an organisation may already possess or are seeking to acquire in context. Figure 6 presents an overview of a conceptual reference architecture that is independent from specific technologies, and may be used to consider where gaps may exist and how they might be filled.

An event correlation engine, event bus, event store and alert store are typically all at the core of traditional SIEM systems. In contrast, one of the items that will typically mark out an analyst-centric solution is the centrality of the graph store. This stores contextual information, key flags/indicators identified by streaming event correlation and filtering, along with linking information that may be derived automatically or by analysts. A graph store can also support automated graph analytics, which are likely to be computationally intensive and therefore will not necessarily be real time.

Similarly, augmentation and lookups on relevant entities – eg, IP addresses, will typically be done asynchronously against the graph store. It will also underpin any graph visualisation.

By way of illustration, Figure 7 below contains some examples of Open Source or free technologies that fit into the areas above that can be readily downloaded and may be useful to fill in gaps. Many proprietary vendors will also provide some or all of these capabilities pre-integrated.

Setting up a full security analytics capability can be an expensive and time-consuming proposition, although increasingly one that is a necessary cost of doing business in many industries.

Many studies are now showing that enterprises are beginning to prioritise their spending on cyber-security, but it can be difficult to put together a cost benefit analysis for security investments: in practice risks and impacts of a breach are often impossible to quantify. This can be countered to some extent by benchmarking against industry peers.

It is also worth noting that it is often possible to use security tools to improve operational monitoring and intelligence, and at times (particularly when consolidating application level data) they can be useful for some types of business intelligence and reporting. Caskey's paper presenting a 'Working Theory of Monitoring' provides a reference architecture with a subset of the components described in the reference architecture above.

Component	Examples
Packet capture and filtering/protocol reconstruction	Suricata, Snort
Log aggregation and structure extractions	Logstash, Flume
Event bus	Apache Kafka, ZeroMQ, RabbitMQ
Event correlation and filtering	Apache Storm, Apache Spark, Esper
Graph store	Neo4j, Tital
Automated graph analytics	Apache Giraph, Gremlin
Event store	ElasticSearch, Apache Accumulo
Metric store	Graphite
Search/ visualisation/ dashboard frontends	Kibana, Grafana, Grphi, Maitego

Figure 7: Examples of Open Source/free components.

Conclusion

Effective defence against sophisticated cyber-attacks is increasingly necessary across a wide range of industries. It will require a change in approach for security teams to use more of the techniques utilised in other security disciplines– based upon a balanced approach of deterrence, denial, degradation and detection of attackers at all points in the kill chain. At the core of these techniques are those associated with intelligence analysis, which can augment existing SIEM tools and empower security analysts to provide flexible and responsive defence for the enterprise.

Organisations should reflect carefully on whether they may be at risk from this type of attack, and if they should step up investment in this area to keep pace with their peers. By evaluating their current capabilities against the provided reference architecture, weak points that should be closed will be identified to allow them to meet the challenges ahead.

Bibliography

Oram, Mark et al. 'Security Intelligence Report Volume 15: January 2013 to June 2013'. Microsoft. Accessed Feb 2014. http://download.microsoft.com/download/5/0/3/50310CCE-8AF5- 4FB4-83E2-03F1DA92F33C/Microsoft_Security_Intelligence_Report_Volume_15_English.pdf.

US Department of Defense. Joint Publication 3-13 Information Operations, February 2006. www.dtic.mil/doctrine/new_pubs/jp3_13.pdf.

Willison, Robert et al. 'Overcoming the insider: reducing employee computer crime through Situational Crime Prevention'. Communications of the ACM, 52(9): 133137, 2009. http://doi.acm.org/10.1145/1562164.1562198.

Selvi, Jose, 'Covert Channels over Social Networks', March 2012. SANS. Accessed Feb 2014. www.sans.org/reading-room/whitepapers/engineering/covert-channels-social-networks-33960.

CrowdStrike Global Threat Report: 2013 Year in Review'. Crowdstrike, Jan 2014. Accessed Feb 2014. www.crowdstrike.com/global-threat/index.html.
'Determined Adversaries and Targeted Attacks: The threat from sophisticated, well-resourced attackers', Microsoft, June 2012. Accessed Feb 2014. www.microsoft.com/en-us/download/details.aspx?id=34793.

Schneier, Bruce. 'Why It's Important to Publish the NSA Programs'. Schneier on Security, October 2013. Accessed Feb 2014. https://www.schneier.com/blog/archives/2013/10/why_its_importa.html.

Charney, Scott, 'Rethinking the Cyber Threat – A Framework and Path Forward'. Microsoft, May 2010. Accessed Feb 2014. www.microsoft.com/en-us/download/details.aspx?displaylang=en&id=747.

Anderson, Ross. 'Security Engineering: A Guide to Building Dependable Distributed Systems (Second Edition)', Wiley, 2008, www.cl.cam.ac.uk/~rja14/book.html.

'Big Data Analytics in Cyber Defence', Ponemon Institute. Accessed Feb 2014. www.ponemon.org/library/big-data-analytics-incyber-defense/.

Willhelm, Wesley Kenneth, 'The Fraud Management Lifecycle Theory: A Holistic Approach to Fraud Management.', Journal of Economic Crime Management, Spring 2004, Volume 2, Issue 2. Accessed Feb 2014. https://library.utica.edu/academic/institutes/ecii/publications/articles/BA309CD2-01B6-DA6B-5F1DD7850BF6EE22.pdf.

Jonas, Jeff; Harper, Jim, 'Effective Counterterrorism and the Limited Role of Predictive Data Mining'. Cato Institute, December 2006. Accessed Feb 2014. www.cato.org/publications/policy-analysis/effectivecounterterrorism-limited-role-predictive-data-mining.

'Intelligence Analysis', Wikipedia. Accessed Feb 2014. http://en.wikipedia.org/wiki/Intelligence_analysis.

Heuer, Richards, 'Psychology of Intelligence Analysis', History Staff, CIA Centre for the Study of Intelligence, 1999. Accessed Feb 2014. www.au.af.mil/au/awc/awcgate/psych-intel/.

Hayes, Ashton, 'Defending Against the Unknown: Antiterrorism and the Terrorist Planning Cycle', The Guardian, April 2008. Accessed Feb 2014. https://www.hsdl.org/?view&did=486223.

Hutchings, Eric et al. 'Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains'. Lockheed Martin. Accessed Feb 2014. www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf.

Reidy, Patrick, 'Combating the Insider Threat at the FBI: Real World Lessons Learned'. Black Hat USA 2013, July 2013. Accessed Feb 2014. http://media.blackhat.com/us-13/US-13-Reidy-Combating-the-Insider-Threat-At-The-FBI-Slides.pdf.

Strand, John; Asadoorian, Paul. 'Offensive Countermeasures: The Art of Active Defence'. Create Space, July 2013. ISBN-10: 1491065966.

Yarochkin, Fyodor et al. 'Hunting in the Shadows: In Depth Analysis of Escalated APT attacks'. Black HatUSA 2013, July 2013. Accessed Feb 2014. https://media.blackhat.com/us-13/US-13-Yarochkin-In-Depth-Analysis-of-Escalated-APT-Attacks-Slides.pdf.

Dickson, Caskey L. 'A Working Theory of Monitoring'. 2013. Accessed Feb 2014. https://www.usenix.org/sites/default/files/conference/protected-files/dickson.pdf.

About BJSS

BJSS is an award-winning delivery-focused IT Consultancy.

With over 20 years' software delivery and IT advisory experience, we are renowned for technical excellence, cost-effective delivery and our proven BJSS Enterprise Agile approach.