SOC 2030: Security Operations Centres are broken, let's fix them

Received (in revised form): 20th June, 2018

Kerry Matre

is the Head of Security Operations Strategy at Palo Alto Networks. She has been defining the steps for customers to transform their architecture and security operations to the next level: increasing prevention, reducing risk and enabling smart people to do smart things within their operations. Her background in security operations has provided insight into over 150 Security Operations Centres (SOCs). Having been involved in assessing the capabilities and effectiveness of SOCs in over 25 countries worldwide, Kerry has a unique view into what has worked in a SOC and what has failed in the past. At HP/HPE ArcSight, Kerry co-authored 'The State of Security Operations Report of Capabilities and Maturity of Cyber Defense Organizations' from 2015 to 2017. She has a BSc in computer science from the University of Colorado at Boulder. Her previous experiences include work at IBM, which involved software development, ethical hacking and creating one of the first and largest data marts for security analytics.

Palo Alto Networks, 3000 Tannery Way, Santa Clara, CA, 95054, USA
E-mail: kmatre@paloaltonetworks.com

Abstract Security Operations Centres (SOCs) are facing many challenges today, including a cyber security skills gap hampering the ability to hire and retain staff, an overabundance of low-fidelity data flowing into the SOC, a broken innovation consumption model and a lacking ability to measure capabilities of a SOC. To overcome these challenges, a fundamental change in the approach to SOCs must be made. The changes necessary to allow a SOC to protect an organisation against successful cyberattacks are not just limited to the SOC itself. They require tight integration with groups aligned with the SOC, including network operations, security engineering, and the lines of business themselves. A prerequisite to this tight integration is a clear mission statement of what service the SOC provides to the business, including what it does and does not do. From there, we can begin to alter the inputs and outputs of a SOC through implementation of a prevention-based architecture and mitigation automation, a new security innovation consumption model and continuous measurement of configuration and operational confidence. This paper will walk through the fundamental changes needed to meet the challenges SOCs face today and move towards the adaptive SOC of the future: SOC 2030.

KEYWORDS: cyber security, security operations centre, cyber defence centre, prevention-based architecture, SOC 2030, SOC metrics, adaptive SOC

INTRODUCTION

Security Operations Centres (SOCs) are struggling to keep up with today's threats. They have evolved over the last four decades to protect against the latest attacks and adversaries, but have repeatedly proven ineffective in protecting organisations against successful cyberattacks. SOCs have their beginning in the 1970s and were mainly a function of military and governments.[1] As the private sector began embracing the opportunities of the Internet, the need for SOCs in the private sector followed. Today, most enterprises worldwide have a SOC in some form.

There are many common issues faced by SOCs. One is the ability to hire and retain qualified staff to manage and operate the

SOC. There is a significant gap between the number of workers with cyber security skills and the staffing needs of organisations. This is fuelled by the growing need for human intervention to handle an increasing number of attacks and the turnover rates driven by burnout as well as ample job opportunities in the industry. A side-effect of this is exceedingly high wages — over 50 per cent higher than the average career and almost 15 per cent higher than a systems administrator.[2]

Another common issue faced by SOCs is an overwhelming amount of data that cannot be effectively processed. The lack of automation results in copious quantities of information being ignored or incorrectly utilised in identifying attacks and unusual network behaviour. The issue can be compounded by low-quality use case.

Both issues contribute to SOCs performing well below recommended levels.[3] This paper will explain what fundamental changes need to be made within organisations to address these issues and provide businesses with the ability to prevent successful cyberattacks as we move towards the year 2030.

MISSION OF A SOC

A well-defined mission and scope are critical to the success of a SOC. The purpose of a SOC is to identify, investigate and mitigate threats. SOCs are often tasked with responsibilities beyond the scope of this mission. These tasks can include engineering, network operations, forensics, incident response, compliance and integrations/development. Although these tasks must be handled by an organisation, they should reside outside of the SOC but integrate tightly with it. This is a separation of duties between the 'watchers' in the SOC and the 'doers' outside of it. When a SOC gets involved in 'doer' tasks, it distracts from its mission and can cause bias and conflict of interest.

SOCs must also operate simultaneously in four distinct modes. These include day-to-day operations, crisis, continuous change and hunting. Day-to-day operations include regular identification of threats, investigation and escalations as required and mitigation with closed-loop prevention of the threats. This also includes typical shift-turnover tasks and keeping up-to-date with the latest threats and investigation techniques. Crisis mode is entered when a widespread vulnerability or attack is identified in the organisation or a large breach is discovered. This mode is unpredictable and hinges on successful communications within the SOC and with external entities. Continuous change recognises the fact that SOCs are fluid. There are always new indicators of compromise (IoCs) to account for, new attack vectors, new tools to be trained on and new processes and procedures to keep up with. The fourth mode — hunting for and identifying unknown threats — is often sidelined due to lack of resources. Some organisations rely on external entities to provide this information; however, the SOC is responsible for finding threats specific to their architecture and business.

It is difficult to be successful in each of these modes of operation, while also adhering to the mission of the SOC. A cleanly defined scope for SOC activities, distinct from other pieces of the business, will enable the SOC to focus on its purpose: to identify, investigate and mitigate threats.

OVERWHELMED SOC ANALYSTS

SOC analysts spend most of their time identifying threats. This is due to an overwhelming amount of data feeding into the SOC and extremely low-quality use cases. A common way for analysts to deal with this glut of data is to de-tune sensors or ignore alerts — both resulting in unmitigated risk to the business. The overwhelmed analysts in the SOC spend little to no time activating a closed loop process to feed

known attacks back into the security controls to prevent future successful attacks. The origins of this challenge are two-fold.

First is the rise of security information and event management (SIEM) systems. These software systems were developed so an organisation could feed security data from disparate technologies into a single place. The SOC is then able to see all security information through a single pane of glass. Additionally, these systems allow for correlation between events and the development of use cases to alert on malicious events. Through poor implementations and low adoption of the advanced features of SIEM, the industry has never realised the full potential of the systems and their use has been reduced to post-breach investigation tools. Most organisations failed to implement solid use cases and the result was an overabundance of raw events and false positives flooding the SOC. Analysts deal with this by detuning sensors and ignore an average of 44 per cent of alerts.[4] A 2017 ESG Research Insights Report echoes this, stating that 54 per cent of security operations organisations ignore events/alerts because they cannot keep up with the volume[5] (see Figure 1). As John Petropoulos from Respond Software writes, 'How crazy is it, that in this day and age we are blinding ourselves just so we can see?'[6]

The second origin of the challenge is the ever-increasing volume and sophistication of attacks being launched at organisations. Available toolsets have become more sophisticated, while the sophistication required of the threat actors to use these toolsets has diminished. Attackers have turned successfully to automation to carry out their breaches. SOCs have attempted to scale with people to combat the threats and

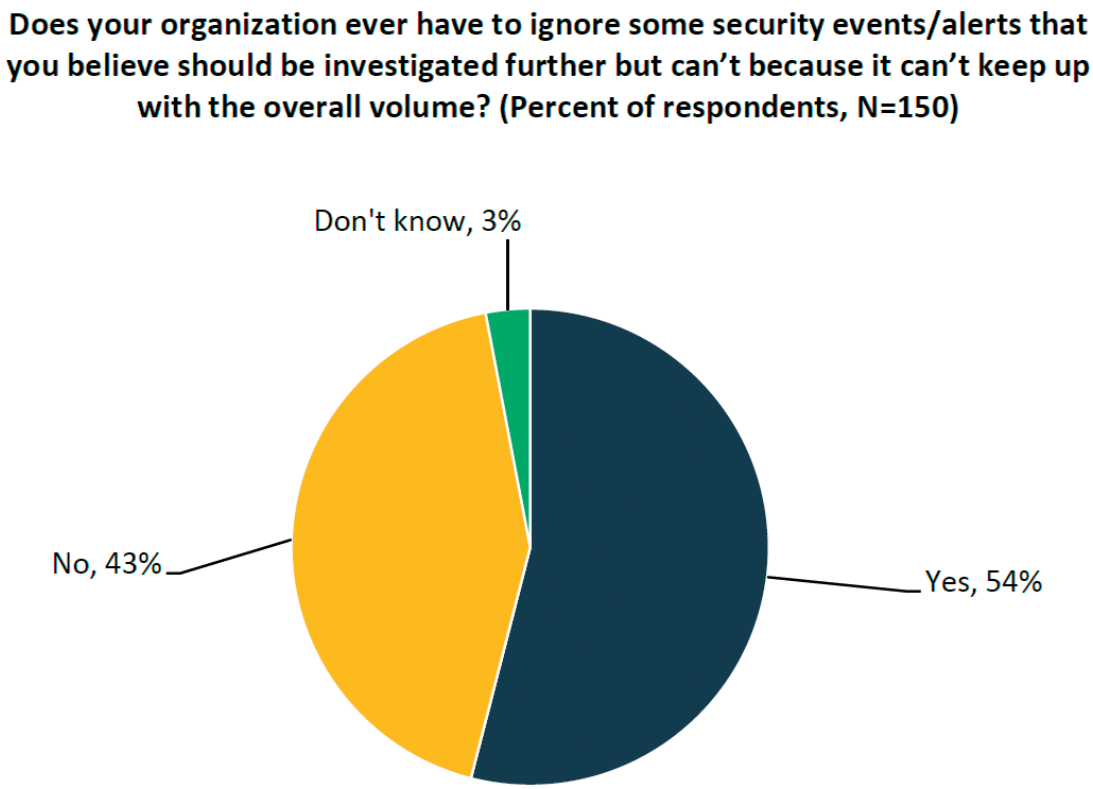

Figure 1: Propensity of organisations to ignore events/alerts due to volume
Source: 2017 ESG Insights Research Report

have failed due to a lack of skilled workers and a finite limit to the ability to match the scale of the attacker's automation (see Figure 2).

Excess amounts of low-fidelity data and false positives flowing into the SOC leaves analysts overwhelmed and ineffective, leading to burnout and high rates of turnover. An innovative approach is required to solve for these challenges.

CHANGING THE INPUTS AND OUTPUTS

Fundamentally changing the inputs and outputs of a SOC will change how a SOC operates. Correctly implementing a prevention-based architecture will lead to less low-fidelity information flowing into the SOC and automated mitigation of threats with implementation of closed-loop security controls will enable blocking of future attacks. A prevention-based architecture consists of four parts.

Consistent protection

Consistent controls across the network, endpoint and cloud will ensure that there are no gaps in protection and will reduce the attack surface. The same use cases should be implemented across the environment for the best ability to prevent attacks or limit the scope of impact. These use cases include credential theft, lateral movement, escalating privileges, data exfiltration and malware infection. There will be some differences per deployment scenario (eg spinning up new VMs in a cloud instance); however, the more consistency applied will result in easier management and increased protection.

Centralised management

Implementing a system for centralised management of security controls will result in faster and more consistent application of security controls. As stated above, consistent security controls will lead to increased protection. Another benefit of centralised management is ease-of-use for the operator; controls can be pushed out to multiple devices with a single click instead of having to access each instance separately. It also allows updates to be rolled out at once so that different areas of protection do not fall out of sync with the rest of the business. An added benefit is a reduction of audit and compliance efforts.

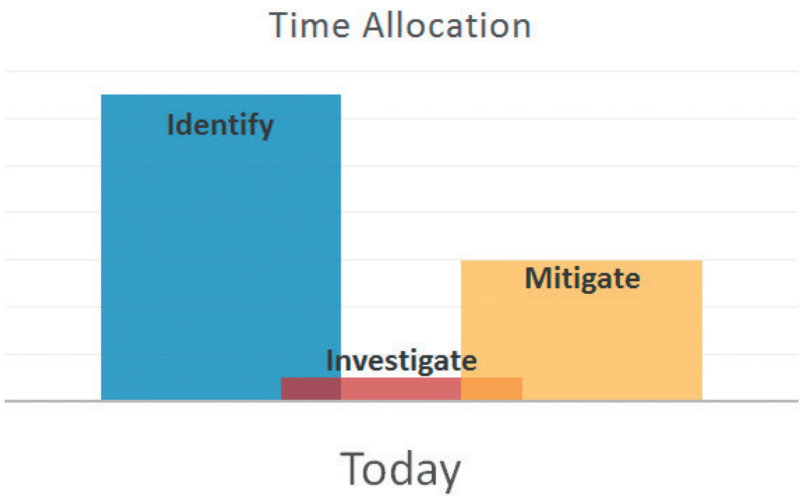

Figure 2: Time allocation of a SOC analyst
Source: Author

Automated threat prevention

Often, organisations include a manual effort to take in IoCs from threat intelligence feeds and manually convert them for security controls to deploy in their environment. Automating this process can shorten this effort considerably and allows security controls to be pushed out in near-real time. The intelligence should come from sources across the globe and allow for a breakdown at the industry and regional levels. Competitors should work together to provide the IoCs to each other for the benefit of all organisations and the detriment of the attackers. Vendors and organisations should then have automation in place to create security controls and have them deployed in their environment. This approach will still require a patient zero — an organisation to experience an attack. However, by sharing the IoCs with other organisations, the scope of impact can be limited.

Prevention based on the business

Traditional controls were based on networks alone — the flow of traffic. These controls were created out of industry best practices and network architecture but did not consider how users interacted with applications and data. The prevention-based approach includes traditional controls in addition to controls based on users, applications and data in the network, datacentre, endpoints and in the cloud. It moves the security controls closer to the business and the preferred targets of attackers. Liaisons between the business and the security organisations will need to be utilised to understand how the business operates and translate that to controls. This results in better protection around the business assets and not just the network.

Each piece of the prevention-based architecture is implemented outside of the SOC with the purpose of increasing the fidelity of information flowing into the SOC and limiting the overabundance of data by blocking attacks before they can propagate through the environment. A prevention-based architecture changes the inputs and the outputs of a SOC, shifting the workload from identifying and mitigating threats to investigating situations and allowing smart people to do smart things. Greg Martin, CEO and co-founder of JASK, states: 'We need to give this work to machines and put our humans into higher-level roles within the SOC.'[7] When machines and AI address the repetitive tasks, managers can up-level the tasks on which analysts are working, leading to higher job satisfaction and better outcomes (see Figure 3).

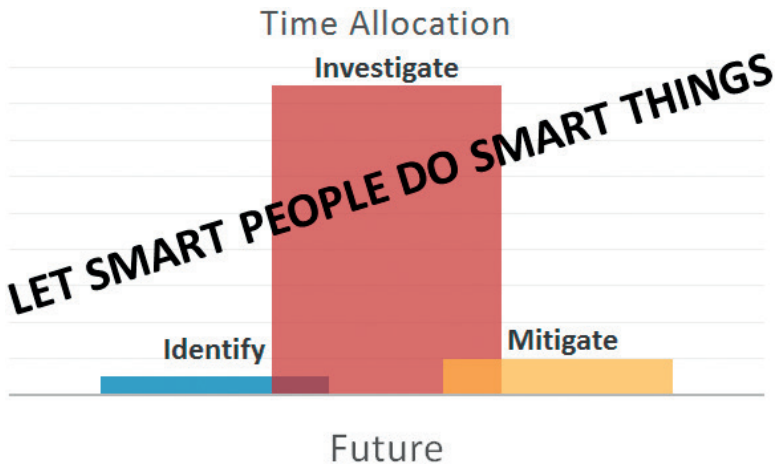

Figure 3: Desired time allocation of a SOC analyst
Source: Author

CONSUMING INNOVATION IN A NEW WAY

Another necessary change surrounds how we consume innovations. By some estimates, SOCs have an average implementation of 40 tools. They are overrun with 'shiny new object' tools and lack a proper investment strategy. 'Companies frequently purchase technology point solutions but fail to bring the data together for effective risk remediation and threat detection.'[8] This overinvestment in single-purpose tools can lead to three additional challenges:

- **CISOs are overrun with vendor-management responsibilities**: Meeting with vendors and negotiating contracts can take up a significant amount of a CISO's time, distracting them from security functions. The focus shifts to tools instead of outcomes for the business.
- **Technologies require an agent to be installed for use**: This proliferation of installed agents has a negative impact on the user experience within the business (eg drained batteries on mobile devices, network slowdowns by the transfer of log data, memory loss due to requirements for operation by each agent).
- **New tools require integration and enablement within the SOC**: Development cycles are needed to integrate each new tool into the technology stack and the processes within a SOC. SOC staff require training on new technologies including new user interfaces and understanding how the use of the tool will change their existing processes and procedures surrounding event identification, investigation and mitigation. Some tools provide benefits to the SOC effectiveness; however, this consuming of innovation is not sustainable.

Additionally, SOCs often use only a small portion of the tools they have in place, leading to sub-par protection and duplicate feature purchases. Research conducted by ESG explains that purchasing too many point tools can be counterproductive because organisations 'don't have adequate time to configure, deploy, and operate new tools to achieve their full potential'.[9]

A shift to a platform approach allows SOCs to consume innovations as features instead of tools, reducing vendor-management requirements. A platform features a main data source on which feature applications can act on that data, allowing the fast-paced innovation of the industry to continue while giving the consumer a straightforward way to access features without having to develop integrations themselves. The main data source is fed by a single agent, consolidating multiple agent on devices, reducing the issues experienced at the device level. A platform approach also enables a centralised enforcement point for security controls to be pushed back out to the devices (see Figure 4).

With the ease of adopting new security innovations, SOCs will be better prepared to handle advanced threats without the distraction of tool implementation and integration. An additional benefit is a reduced barrier to entry for small companies with great innovations that can be applied to security data but lack the resources for agent deployment.

METRICS AND ACCOUNTABILITY

Typically, relevant metrics around the SOC are poorly defined and implemented. They are often inherited from network operations centres (NOCs) and irrelevant to a SOC practice. An example is mean time to resolve (MTTR) or respond. This metric drives the wrong behaviour in a SOC. It focuses on the speed of closing issues instead of full understanding of threats and execution of closed-loop mitigation to prevent future successful attacks. Justin Bajko, VP of customer experience at Expel, states: 'Requiring your analysts to complete an investigation in 10 minutes sounds fine on

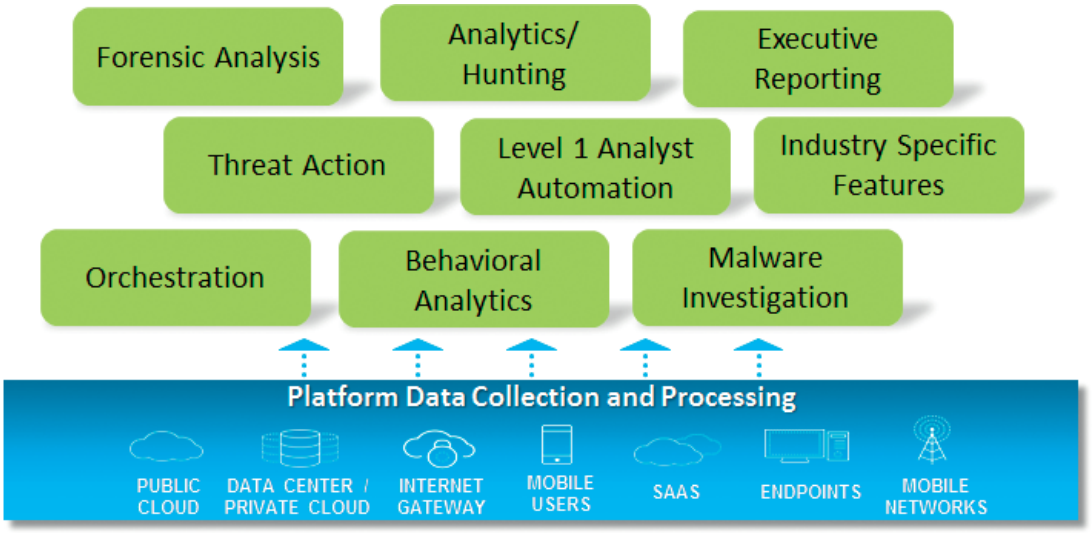

Figure 4: Platform data collection and processing with applications that can act on the data
Source: Author

the surface, but if you constrain an analyst's ability to actually dig into an incident and find out what's really going on for the sake of time, you're likely to miss critical details and negatively impact your response to that incident.'[10]

Another poor metric is based on the number of feeds flowing into a SOC. This is often used as a gauge of project completion (especially in SIEM implementations). Organisations will claim success when ten (for example) data feeds are flowing into their SIEM. However, this does not reflect capabilities of a SOC. Instead, SOCs should focus on the number of use cases implemented with those data feeds. Each data feed should be associated with a use case and outcome, otherwise it is not of use in a SOC and just creates noise for an analyst.

In general, security metrics should fit into two categories and should be continuously measured:

Configuration confidence: Configuration confidence is knowing that an environment is operating as intended and configured following best practices. A business should know what percentage of their environment has applied security controls in place. Any gaps should be understood by the SOC. The SOC should also have instant notification if any of the security controls are not running or misconfigured. If any of the controls are not configured in alignment with best practices, the business should be able to identify the discrepancies and understand why the deviation exists.

Example metrics:

- Enabled controls — a high-level view of controls enabled by business unit, geographical location or even device group can pinpoint gaps in security controls;
- Disabled controls — this metric will indicate controls that were enabled but are no longer running. Investigation into why this change occurred can follow;
- Configuration drift — discrepancies between baseline configurations and current configuration templates can indicate the need for a new configuration baseline or a deviation from intended and effective controls;
- Best-practice alignment — measurement of best-practice alignment can uncover weaknesses in controls. Ideally, this is a heatmap visualisation (see Figure 5) that

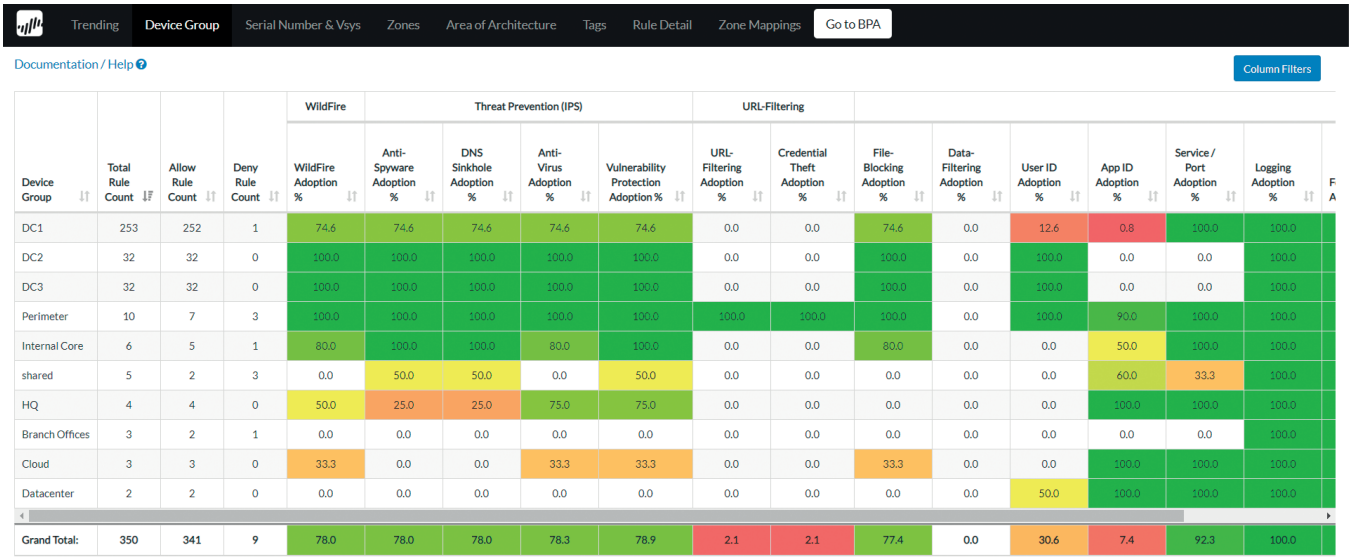

Device Group	Total Rule Count	Allow Rule Count	Deny Rule Count	WildFire	Threat Prevention (IPS)				URL-Filtering								
				WildFire Adoption %	Anti-Spyware Adoption %	DNS Sinkhole Adoption %	Anti-Virus Adoption %	Vulnerability Protection Adoption %	URL-Filtering Adoption %	Credential Theft Adoption %	File-Blocking Adoption %	Data-Filtering Adoption %	User ID Adoption %	App ID Adoption %	Service / Port Adoption %	Logging Adoption %	F A
DC1	253	252	1	74.6	74.6	74.6	74.6	74.6	0.0	0.0	74.6	0.0	12.6	0.8	100.0	100.0	
DC2	32	32	0	100.0	100.0	100.0	100.0	100.0	0.0	0.0	100.0	0.0	100.0	0.0	0.0	100.0	
DC3	32	32	0	100.0	100.0	100.0	100.0	100.0	0.0	0.0	100.0	0.0	100.0	0.0	0.0	100.0	
Perimeter	10	7	3	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	0.0	100.0	90.0	100.0	100.0	
Internal Core	6	5	1	80.0	100.0	100.0	80.0	100.0	0.0	0.0	80.0	0.0	0.0	50.0	100.0	100.0	
shared	5	2	3	0.0	50.0	50.0	0.0	50.0	0.0	0.0	0.0	0.0	0.0	60.0	33.3	100.0	
HQ	4	4	0	50.0	25.0	25.0	75.0	75.0	0.0	0.0	0.0	0.0	0.0	100.0	100.0	100.0	
Branch Offices	3	2	1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0	
Cloud	3	3	0	33.3	0.0	0.0	33.3	33.3	0.0	0.0	33.3	0.0	0.0	100.0	100.0	100.0	
Datacenter	2	2	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	50.0	100.0	100.0	100.0	
Grand Total:	350	341	9	78.0	78.0	78.0	78.3	78.9	2.1	2.1	77.4	0.0	30.6	7.4	92.3	100.0	

Figure 5: Example of heatmap-style report of enabled controls by device group
Source: Palo Alto Networks Best Practice Assessment tools[11]

you can drill down into to discover what is misaligned and how to make necessary updates.

Operational confidence: Operational confidence is knowing that processes are in place around the security controls to meet the business needs. The business wants to know: What is the threat? Am I affected? What is the risk to the business? What is the plan to fix it? The SOC should have metrics in place to ensure that these answers can be provided when needed. This includes knowing that the SOC has the capacity to handle incoming events, which requires ensuring that the controls in place are operating as intended.

Example metrics:

- Events per analyst hour (EPAH) — if this number is too high, the analyst becomes overwhelmed. The effect is that analysts miss incidents or rush through investigations — not fully addressing threats. The resulting outcome: ineffective analysts and risk exposure to the organisation;
- Duplicate incidents — this designates a breakdown in the closed-loop process of incident mitigation feeding back in to security controls. If proper security controls are in place, a SOC should never see the same incident twice;
- Alerts for known threats — this indicates a security controls failure. If a threat is known, security controls to block the threat should already be in position;
- Deviation from SOC procedures — any deviation from a SOC procedure should be reported. If the same deviation occurs multiple times, then either a procedure needs to be updated or additional analyst training is required.

The responsibility of collecting most metrics should not fall upon organisations. Vendors should be held accountable for proving measurements of effectiveness of their tools and the adoption of best practices. This provides the business with an understanding of any gaps in its security controls and operations, as well as a roadmap for improving the protection provided to the business. The effort to improve the security controls will require strict process integration with other functions in the business (engineering, development, lines of business, etc.). In many organisations this tight

coordination does not exist and will need to be established and driven from the executive level. Effective business-level metrics can drive this sustainable business support and focus.[12]

OUTCOMES

Any business change should be tied to outcomes. This is true for any change to security operations as well. As we move towards SOC 2030, each piece of the transition to an adaptive SOC should provide outcomes directly tied to the enablement of the business. The outcomes should be defined before changes are made so they can be measured before and after the change. This not only provides verification of effect on the business but also improves the SOC brand and leads to better cohesion between the SOC and other areas of the business.

Benefits of this approach have been observed at individual organisations; however, no industry study has yet been performed to show the average gains to be expected from adopting various pieces of this approach. Outcomes observed in individual organisations include:[13]

Consistent protection:

- 90 per cent reduction in support tickets to reimage machines from malware infections;
- Reduced threat alerts from ~20m per day to ~1.2m per day — nearly 80 per cent reduction in noise;
- Automatic blocking of 95 per cent of events;
- Stabilised SOC headcount;
- Reduced EPAH from 200 to 20.

Centralised management:

- 40 per cent gain in administration efficiency.

Automate threat prevention:

- Implementation of security controls based on IoCs reduced from days to minutes due to automated threat prevention;
- Threat intelligence processing reduced by 50 per cent from automation;
- EDLs' turned threat response to blocks within 5 minutes.

Prevention based on the business:

- Investment of a business liaison instead of a security analyst resulted in additional inclusion in the process of the new development tool purchasing process. Allowed for security controls around the application to be put in place before tool go-live.

Alternative innovation consumption:

- Time to adopt new security visibility feature reduced from weeks to 15 minutes due to a SaaS model consumption of the feature;
- 28 per cent reduction in security spend resulting from adopting a platform approach.

Meaningful metrics:

- Significantly increased feature adoption;
- EPAH metrics drove automation initiative to reduce noise in the SOC.

Results experienced will vary by organisation and deployment scenarios. However, the numbers above demonstrate a meaningful change in the number of threats being blocked — reducing the amount of data flowing into a SOC, thereby allowing a SOC to focus more on investigation and let automation take care of identifying and mitigating threats.

One observed SOC adopted the pieces of this new, adaptive approach, including a prevention-based architecture, automated security control updates driven by the latest threat intelligence and metrics based on

configuration and operational confidence. As a result, they processed 2.9bn raw log events in 2017, resulting in 3,900 events requiring various levels of investigation (see Figure 6). Automation, including the gathering of contextual data and orchestration realised through SOC playbooks, assisted first line analysts in managing the context-rich events, resulting in 15 true incidents, eight of which were escalated. Fifteen incidents is a sustainable volume, requiring full incident investigations and response for a SOC. They attribute this achievement to the prevention of attacks by controls upstream from the SOC, substantial automatic remediation of alerts and the closed-loop processes implemented to update security controls to prevent repeat alerts.[14]

CONCLUSION

To overcome many challenges that SOCs are facing today, a fundamental change in the approach to SOCs must be made. This change will not happen overnight but needs to begin today and continue to progress over the next decade. SOCs of the future will be adaptive: adaptive to threats, adaptive to new technologies, adaptive to available skill sets of people, and adaptive to the needs of the business. The changes discussed in this paper that support the transition to an adaptive SOC include:

- Altering the inputs and outputs of a SOC through implementation of a prevention-based architecture and mitigation automation;
- Consumption of emerging security innovations via a new model;
- Continuous measurement of configuration and operational confidence.

Implementing a prevention-based architecture will fundamentally change the inputs and outputs of a SOC. It will reduce the amount of noise and overwhelming event volumes by blocking attacks early in the attack life cycle. Security controls are consistent throughout the environment — from cloud to the network to the endpoint — and are centrally managed to reduce administration time. Threat intelligence is automatically processed with this architecture and generates automatic controls or suggestions for controls to be implemented. This reduces the turn-around time from IOCs to security controls from days to minutes. Additionally, controls are placed as close to the business as possible. This means pushing controls out to the new perimeter of enterprises which are the users, applications and data.

Adopting a platform approach to security controls will allow for the consumption of innovations while minimising the unused toolsets as we see today. Integrations of toolsets should be handled by the vendors, freeing up security resources. This approach allows the SOC to be adaptive to innovations without having to install agents and disrupt users. With this approach, enterprises can pick and choose which combination of tools will work best to meet their business needs. It will also drive towards a SaaS and Cloud-enabled model which allows for rapid modifications to be made to technology and processes to combat the advancing threats.

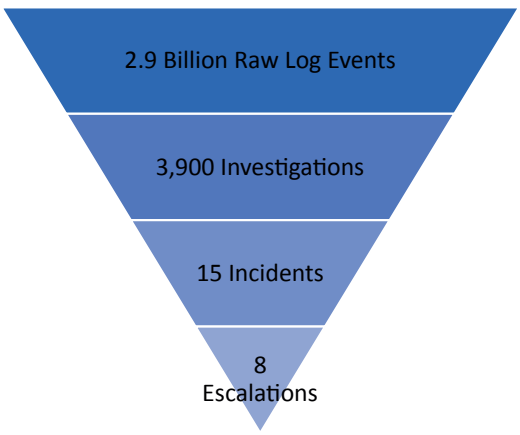

Figure 6: Reduction of low-fidelity events to incident escalations of a large enterprise
Source: Author

Finally, you cannot manage what you cannot measure. Employing a solid set of metrics will provide confidence to the business that it can handle a cyberattack. Metrics providing configuration confidence show the level of security controls implemented and configured according to best practices highlighting gaps in the controls. Operational confidence is gained by establishing metrics related to the operation of the environment including process execution. These metrics provide visibility into the ability of a security organisation to operate effectively during the standard modes of a SOC: day-to-day operations, continuous change, crisis and hunting. Measuring this confidence during the transformation to an adaptive SOC model will also substantiate progress and prevention capabilities.

Threats will evolve drastically as we advance towards 2030, and SOCs must become adaptive to face new risks. Reducing clutter in the form of low-fidelity data, unused toolsets and misleading metrics is necessary. Organisations must seize the opportunity to adopt this innovative approach and begin to implement these changes with the ultimate purpose of preventing successful cyberattacks.

References

1. HP ESP Security Intelligence and Operations Consulting Services (2013), '5G/SOC: SOC Generations', available at http://www.cnmeonline.com/myresources/hpe/docs/HP_ArcSight_WhitePapers_5GSOC_SOC_Generations.PDF (accessed 30th July, 2018).
2. Bureau of Labor Statistics (2017), 'May 2017 National Occupational Employment and Wage Estimates United Stated', available at https://www.bls.gov/oes/current/oes_nat.htm (accessed 30th July, 2018).
3. Micro Focus (2018), '2018 State of Security Operations: Report of the Capabilities and Maturity of Cyber Defense Organizations Worldwide', available at https://software.microfocus.com/es-es/assets/enterprise-security-products/state-of-security-operations-2018-report (accessed 30th July, 2018).
4. CISCO (2017), '2017 Security Capabilities Benchmark Study', available at http://blog.executivebiz.com/2017/02/cisco-study-44-of-daily-security-alerts-are-not-investigated/ (accessed 30th July, 2018).
5. ESG Research (2017), '2017: Security Operations Challenges, Priorities, and Strategies' Insights Report, available at https://resources.siemplify.co/hubfs/PDF%20Downloads/ESG-Research-Insights-Report-Siemplify.pdf?t=1528489300042 (accessed 30th July, 2018).
6. Petropoulos, J. (2017), 'High Volume Event Feeds', Respond Software, available at https://respond-software.com/blogs/blogs/high-volume-event-feeds-1 (accessed 30th July, 2018).
7. Zeichick, A. (June 2018), 'Reinventing and scaling the SOC with AI: Helping humans, not replacing them', Telecom Times, available at https://telecomtimes.com.au/2018/06/08/reinventing-and-scaling-the-soc-with-ai-helping-humans-not-replacing-them/ (accessed 30th July, 2018).
8. *Ibid.*, note 3.
9. *Ibid.*, note 5.
10. Justin Bajko, J. (2017), 'Mistakes to avoid when measuring SOC performance' Expel, available at https://expel.io/wp-content/uploads/2017/11/WP-measuring-soc-performance.pdf (accessed 30th July, 2018).
11. See https://www.paloaltonetworks.com/documentation/80/best-practices/best-practices-data-center/data-center-best-practice-security-policy/use-palo-alto-networks-assessment-and-review-tools (accessed 13th August, 2018).
12. *Ibid.*, note 3.
13. Statistics gathered through on-the-job observations and measurements.
14. Statistics gathered from the 2017 dataset of a large Americas-base technology enterprise.