Measuring Security Effectiveness:

A Critical Requirement for Cybersecurity Leaders

By Matt Hartley – ISSA Senior Member, Northern Virginia Chapter

Security leaders continuously look for ways to improve and evolve their security strategy. One real challenge they face is a lack of approaches for measuring the efficacy of their security program. This can be overcome with an emerging technology frequently categorized as breach and attack simulation, although when examined more broadly is better characterized as security validation.

For many years, cybersecurity leaders have struggled to accurately assess their security posture against real-world threats. They invested heavily in security technologies to prevent new and evolving attacks, and in many cases these technologies were installed and left running with little more than a gut feel of the value those investments offered them. In the best cases, leaders employed human penetration tests and red team assessments on their organizations to discover weaknesses in their security posture. As boards and executive leadership demanded an accounting of the value of these investments, many pointed to raw volumes of detections, the headlines and dollar and employment impacts of breaches in similar organizations, and more recently, to the ever increasing sophistication of the threat environment as justification for their expenditures.

Forward-thinking leaders are now looking for ways to manage cyber like any other enterprise business unit–with metrics-based management. Today, cybersecurity leaders can leverage a new capability for their toolbox, categorized by some as breach and attack simulation, although when viewed more broadly can be better characterized as security validation technologies. These technologies automate a spectrum of testing and preparedness capabilities and even support business outcomes by capturing quantitative measures of effectiveness. These can be applied to strengthening and streamlining security programs and justifying investments as part of a strategic business framework.

Background

Understanding this space and its potential future requires a brief review of its past. For many years, security teams invested in countless defensive technologies without being able to accurately determine how effective they were in isolation or inside their larger security architecture. In many cases, they relied on anecdotal evidence; in other cases, their strategy amounted to redundancy, adding more technology when they discovered weaknesses in existing ones. These teams relied on "point-in-time assessments that require them to "cobble together" data from disparate systems to truly under-

stand the organization's security posture."[1] These included human penetration testing, "a test methodology intended to circumvent the security function of a system," and the outcomes of red teams who "emulate[d] a potential adversary's attack or exploitation capabilities against an enterprise's security posture," to identify weaknesses and to see how well their defenses protected their organization from determined attackers.[2] However, the lack of experienced talent made these tests difficult to conduct and expensive to acquire. Numerous technologies were developed to verify preparedness against threats actors by performing adversary simulation, "a method to test a network's resilience against an advanced attacker."[3] Security practitioners began to explore broader automation to more easily and reliably test their defenses.

A new wave of technologies led Gartner in 2017 to make the security industry aware of what they called "breach and attack simulation" (BAS) technologies, which provided security leaders the means to test and improve their security posture. While human operators and manual approaches still provided a valuable means to test security postures, in Gartner's view, providing "continuous testing…is the key advantage of BAS technologies [to] validate that security infrastructure, configuration settings, and prevention technologies are operating as intended."[4] As they and others saw it, "security testing is so challenging for technical professionals focused on security operations that many don't try it." Their research highlighted the value BAS offered its users in the form of easy-to-use automation to test and harden the defenses of organizations.

These technologies were presented as covering the gamut of security controls testing, results, mitigation prioritization, and even process testing. In reality, they did offer new automated, repeatable ways to launch simulated and real attacks on a general target address, on specific controls, or on an entire network presence of an organization in order to assess the actual performance of one or more security controls versus what is expected or if an attacker could gain access to the organization and move laterally throughout its network presence. Some provided results in simple stoplight dashboards displaying red for compromised or ineffective and green for protected or effective; others provided numeric measures like the number or percentage of attacks detected and blocked. The variety of technology options presented was broad, which the authors attempted to address in part by calling attention to one user's differentiation of manual red teams and the use of these BAS technologies as "penetration testing helps answer the question 'can they get in?'; BAS tools answer the question 'does my security work?'" While the authors acknowledge this view is not adequately refined, their attempt to differentiate BAS as a single area falls short.[5]

Illuminating these technologies led to a positive outcome—security professionals became aware of automation technologies that enabled them to actually validate and quickly improve their security defenses. However, Gartner's view that these technologies fall into a single category has somewhat hindered a more comprehensive understanding of the out-

1 "Forrester Study Highlights 'a False State of Confidence' When It Comes to Enterprise Cybersecurity," Continuity Central, Sep 27, 2019 - https://www.continuitycentral.com/index.php/news/technology/4462-forrester-study-highlights-a-false-state-of-confidence-when-it-comes-to-enterprise-cyber-security.

2 See "Penetration Testing," US National Institute of Standards and Technology (NIST) Computer Security Resource Center - https://csrc.nist.gov/glossary/term/penetration-testing and "Red Team," NIST Computer Security Resource Center - https://csrc.nist.gov/glossary/term/red-team.

3 "List of Adversary Emulation Tools," PenTestIT - https://pentestit.com/adversary-emulation-tools-list/.

4 Greg Young, "Hype Cycle for Threat-Facing Technologies, 2017," Gartner, Jul 17, 2017 - https://www.gartner.com/en/documents/3762274.

5 Anton Chuvakin & Augusto Barros, "Utilizing Breach and Attack Simulation Tools to Test and Improve Security," Gartner, May 17, 2018 - https://www.gartner.com/en/documents/3875421.

comes enabled by technologies in this space. Looking briefly at the history and approach within each discipline denoted in each quarter of figure 1, we can identify an important and foundational key to understanding this space: penetration testing and controls assessments were typically oriented internally while red teams and threat preparedness were usually oriented externally. Given that in each case these activities were largely manual at first, and technologies were developed to automate actions and make teams more efficient specific to each area, they continued to focus inwardly or outwardly. This specificity is critically important—many technologies in the BAS space are really focused on a single discipline, like technologies to automate red teaming attacks to compromise systems and networks, or technologies to automate the assessment of individual security controls. BAS does not easily allow for an overarching fusion of some or all of these disciplines into a single technology. Figure 1 in total depicts this higher-order view of reality, that each approach is necessary for a strong security posture and that the intersection covers the totality of approaches for strong security. These coalesce holistically into a technology area that can be called security program validation.

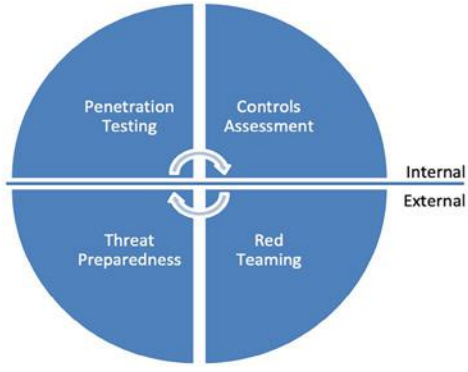

Figure 1 – Security program effectiveness

These validation technologies are now evolving and others will emerge that combine automation for all these disciplines inside a single platform in order to prove broader security program effectiveness. Security program validation technologies offer the most potential, given their ability to deliver across the widest variety of use cases. At the most fundamental level, these technologies will automate and support security control assessments, purple teaming, risk management, and operationalizing threat intelligence, each of which will be explored in more detail below. As these technologies are further deployed and evolve new capabilities, one can easily project the possibility of these systems providing quantitative data from each approach to support business outcomes like calculating return on investment within a security program. While not an exhaustive list, these use cases demonstrate what a true security validation platform could offer to improve security programs.

Continuous control assessments

The fundamental value of these technologies is their ability to automate the assessment of controls in an "efficient and consistent way to measure the effectiveness of existing security detection capabilities and operations."[6] In essence, they automate the process of executing a series of attacks across production infrastructure, correlating and collecting the results of each attack from logs and security tools, and summarizing the overall results across all attacks and controls. The attacks are typically able to be run individually, in sequence, or as part of a larger scenario. Results are collected in a variety of ways, typically starting with inspecting logs in SIEMs or analytic platforms and sometimes directly from security controls in order to verify if the security control detected and prevented each attack. The overall outcomes are collected and correlated and presented to the user in various forms, like a total count of positive results versus total attacks and the derived percentage efficacy across a variety of viewpoints, such as for each technology across all attacks.

For example, if a sequence of 200 attacks were executed and 122 were detected and 62 were blocked, the detection effectiveness of the controls applicable to that attack would be 61 percent and the prevention efficacy would be 31 percent. The results could be organized by control in order to compare which controls were more capable of defending the attacks, or organized in other ways to look across the larger security posture of the organization. Teams can then focus on specific controls that need improvement, taking actions like confirming and reconfiguring settings and checking for and installing updates as needed. The same series of attacks could then be executed again to validate any improvements in effectiveness. As represented in figure 2, a team could iterate repeatedly between assessment and improvement with the same attacks (as well as others as applicable) to maximize the control's ability to detect and prevent the attacks.

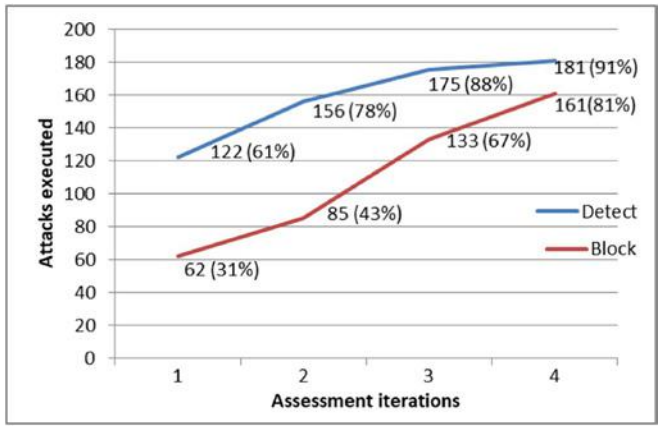

Figure 2 – Graphically represent control assessment progress

The ability to run the exact same test repeatedly, either manually or continuously at some time interval, is a key feature of control assessment technologies. When these technologies first emerged, this was a major improvement over earlier manual penetration testing and red teaming where real repeatability was challenging. The continuous nature of se-

6 Ashley Arbuckle, "Fact vs Fiction: The Truth about Breach and Attack Simulation Tools," SecurityWeek, Jul 25, 2019 - https://www.securityweek.com/fact-vs-fiction-truth-about-breach-and-attack-simulation-tools.

The What, Why, and How of Cybersecurity Asset Management

AXONIUS

By Lenny Zeltser, CISO, Axonius

How can we secure an IT resource if we don't know that it exists or if we don't have visibility into its state? To quote respected industry practitioner Adrian Sanabria, "most security and IT problems begin with visibility." Security practitioners crave visibility into the state of laptops, devices, virtual machines, applications, and users in their organization. Overseeing security aspects of the configuration of such resources is the practice of ***cybersecurity asset management****.*

What Does Cybersecurity Asset Management Involve?

To address security issues, you must discover the gaps, and to do that you need a comprehensive and reliable inventory of your asset. Therefore, cybersecurity asset management involves:

- Obtaining and continually updating an accurate inventory of all IT resources.
- Discover security gaps related to the asset's presence or configuration.
- Enforcing security requirements to rapidly address the identified gaps.

Asset management plays such a foundational role in a cybersecurity program, that CIS Critical Controls lists the need to inventory and control hardware and software assets as its first two security measures. Along these lines, asset management is the first category in the NIST Cybersecurity Framework. For yet another example, consider guidance by the Security and Exchange Commission, which highlights the need to inventory hardware and software so the organization knows where its assets "are located, and how they are protected."

Unfortunately, implementing this process in a reliable, timely, and efficient manner has been one of our industry's major challenges.

Why Don't We All Have Asset Management Already?

If asset management is so important for cybersecurity, why haven't all enterprises implemented it yet?

"Basics are hard," as Adrian Sanabria put it.

In cybersecurity, we're often attracted to exciting-sounding disciplines, say threat hunting or red-teaming. We're drawn to sexy technologies such as machine learning for malware or anomaly detection. We struggle to take a step back to build a foundation for the security program, even if we know it'll enable cool efforts such as spotting intrusions and fighting malware.

Another reason why asset management has been a challenge is the lack of effective tooling. Keeping track of IT resources is often a manual, error-prone process that consumes much time and yields few benefits. For asset management to deliver its full potential, it needs to be automated and easy to implement.

How to Approach Cybersecurity Asset Management?

Here's the good news. Today's enterprises already have many IT and security systems that know about some portion of the organization's assets. The challenge from the perspective of asset management is that these systems typically exist as data silos, requiring cumbersome efforts to get a unified and actionable view on asset details across multiple systems.

Organizations can advance their asset management program by extracting useful configuration and other state data out of these systems. The next step is to clean the data to find useful information across the multiple data sources.

Continue reading on how to solve the top cybersecurity asset management challenges by accessing our latest white paper at **axonius.com/challenges**

curity validation technologies allowed for new ways to detect changes in control configurations and even in their surrounding infrastructure, given that "even if security controls are working today, any change or update to the environment could potentially change this."[7] In other words, it wouldn't be uncommon to discover controls that previously detected or blocked attacks no longer doing so, but finding these unexpected changes were difficult outside a sporadic manual assessment. In contrast, with automated assessments security teams can gain confidence that their defenses continue to operate as expected to prevent adversaries from breaching the organization's systems and networks.

Automated purple teaming

Purple teams emerged in recent years to bridge the experience of a red team assessment with training for the defending blue team personnel. A purple team assessment commonly includes the red team and blue team sitting "side by side to collaborate and truly understand outcomes" achieved.[8] The red team leverages predetermined attacks against specific controls that the blue team can monitor. Purple team methodologies ensure the blue team learns to detect and defeat these attacks in the future through a combination of process improvements and control and infrastructure changes. Purple teams are also valuable for improving the skills of defenders as they dive deeper to learn how attackers think and how to defeat them.

Security validation technologies are a strong choice for delivering purple team engagements. As highlighted in the previous section, these technologies innately offer the ability to select and run attacks safely inside production environments. They highlight the effectiveness of specific controls and offer defending security teams the ability to understand where the controls are falling down. The advantages of using validation technologies for purple teams are many, and include productivity gains in the purple team execution from the inherent technology automation. Another advantage is the repeatability they offer, the ability to run the same tests again to ensure that the combination of process and technology improvements and the lessons the blue team learned are realized for the long-term. These platforms can provide a unified interface through which both red and blue teams can collaborate versus just interacting in person. They also offer the means for blue teams to test themselves against emerging attacks from the global threat landscape, a critical advantage for the security team to proactively explore its readiness for the latest threats.

Operationalizing threat intelligence

Over a decade ago, threat intelligence was largely focused on technical aspects of threats, such as malware, botnets, and spam, represented in these cases by malicious files and hashes, IPs and domains, and spoofed senders. The initial value of this information for many in the security community was to improve their operations and tools, typically as technical context was integrated into security tools to gain better awareness of attacks and improve prioritization. Over time researchers began to center on the threat actors conducting these attacks, tying together the technical characteristics with tactics within a linear attack progression that included other aspects of an attacker's actions, such as what they did before and after an attack.[9] These attacks could be tied together into campaigns being executed around the world across broad geographies and even entire industries. Security teams could then attribute attacks in their environment to these campaigns and the threat actors executing them. This ultimately enabled security teams to learn which actors they faced regularly and allowed them to prioritize proactively preparing to defend against those actors and their evolving tactics.

The MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) framework has emerged as a key resource for security teams attempting this process of defending against the tactics, techniques, and procedures (TTPs) of

7 Adrian Sanabria, "A Primer on Breach and Attack Simulations," MIS Training Institute, Jun 26, 2018 - https://misti.com/infosec-insider/a-primer-on-breach-and-attack-simulations.

8 Joseph R. Salazar, "The Rise of 'Purple Teaming'," Dark Reading, Jun 13, 2019 - https://www.darkreading.com/threat-intelligence/the-rise-of-purple-teaming/a/d-id/1334909.

9 Matt Hartley, "Think Like Your Adversary: Leveraging the Cyber Threat Kill Chain," ISSA Journal (Nov 2014): 20-24 - https://www.members.issa.org/resource/resmgr/journalpdfs/feature1114.pdf.

threat actors. The framework attempts to "relate behaviors to defenses" in ways that are "applicable to real environments," breaking attacks down into their component techniques that can be used to construct scenarios for security "teams to plan [attack] events and for the detection team to verify their progress."[10] To that end, security validation technologies offer new means to operationalize threat intelligence. Rather than only focus on security operations visibility and prioritization, these technologies can leverage threat intelligence to also validate the effectiveness of security controls by TTPs from a variety of attackers as well as by all the TTPs of one or more specific adversaries. In other words, security teams can leverage ATT&CK to perform gap assessments on their defenses to discover what needs hardening.

For example, a security team could thoroughly inspect their environment for the potential of data leakage by executing various techniques that attackers would use to extract data from their environment, leveraging tests aligned to the techniques associated to the MITRE ATT&CK "Exfiltration" tactic.[11] Similarly, a security team might want to ensure they are protected against a specific adversary, such as a nation-state actor that frequently targets organizations in their industry, and would run tests aligned directly to the tactics and techniques of that actor. As a key tenant of security validation, these approaches provide direct evidence of which controls succeeded or failed at detecting and blocking highly-relevant attacks and highlighting corrective actions that teams need to prioritize to strengthen their defenses. Security validation also identifies the adversaries to which an organization is potentially most vulnerable, allowing the security team to prioritize gathering intelligence on those attackers to proactively track and prepare for how they are evolving.

Effectiveness in risk management

Security validation technologies are also valuable within fundamental risk frameworks. For example, the NIST Cyber Risk Management Framework (800-37) includes an "Assess step to determine if the controls selected for implementation are implemented correctly, operating as intended, and producing the desired outcome."[12] While this covers more than just security technologies, the reference in this step to other NIST documents like their report on automated security control assessments clearly indicates the opportunity for security validation technologies to play a key part for organizations conducting risk assessments. These documents outline processes that are very similar to automated control assessments, are mapped to attack models that align to the aforementioned kill chain approaches, and highlight the need to examine not just individual controls but an entire organization's "defense in depth."[13]

In fact, where in many cases risk frameworks use common qualitative measures like high, moderate, low, and none, validation technologies provide numeric effectiveness values that are valuable for more accurate impact analyses and remediation prioritization. And in many cases even values used within frameworks that are hypothetically more quantitative in approach are frequently estimates. For example, measures for control strength and vulnerability against certain threat impacts are commonly determined by fitting a qualitative estimate into one of a few predefined ranges and then obtaining a specific, estimated quantitative value previously associated to that range. As rapid, broader adoption of the Factor Analysis of Information Risk (FAIR) model has occurred, security professionals have not widely recognized the value that control assessment and security validation technologies provide the means to more directly measure the ability of controls to stop specific attacks.[14]

Quantified effectiveness and business metrics

One of the more challenging areas security professionals with a technology-oriented background face is demonstrating their ability to navigate business conversations with their leaders and their boards about the value of the investments they are making. Over the past two decades, these conversations evolved from a technical and personnel focus to the depth of investment and breadth of the security team program. Now, however, leadership demands security be presented within larger organizational governance with real financial accountability.[15] Providing these insights has been a challenge for security leaders, with many professionals taking a largely qualitative approach focusing on discrete technical aspects of their programs. Security validation technologies can fill this void and offer quantitative approaches to meet leadership's requirements, resulting in more effective communication and garnering understanding and trust from executives and boards.

The results of the use cases described earlier can be directly leveraged in business conversations. For example, using graphics like figure 2 are a straightforward means to represent a team's progress improving security posture. Leaders can go even further, using controls assessment results and borrowing from basic business calculations like efficiency,

10 See "MITRE ATT&CK," MITRE, accessed Feb 2, 2020 - https://attack.mitre.org/ and Blake Strom, "Why ATT&CK Was Created," MITRE ATT&CK Blog, Sep 20, 2018 - https://medium.com/mitre-attack/att-ck-101-17074d3bc62.

11 "Exfiltration," MITRE, accessed Feb 2, 2020 - https://attack.mitre.org/tactics/TA0010/.

12 "Risk Management Framework for Information Systems and Organizations," US National Institute of Standards and Technology (NIST), Dec 2018 - https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-37r2.pdf.

13 Kelly Dempsey, et al., "Automation Support for Security Control Assessments," US National Institute of Standards and Technology (NIST), Jun 2017 - https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8011-1.pdf.

14 See "FAIR Risk Management," FAIR Institute - https://www.fairinstitute.org/fair-risk-management and Christina Dulovich, "4 Rules for a Successful Quantitative Cyber Risk Analysis," FAIR Institute Blog, Feb 6, 2020 - https://www.fairinstitute.org/blog/4-rules-for-a- For example, measures for control strength and vulnerability against certain threat impacts are commonly determined by fitting a qualitative estimate into one of a few predefined ranges and then obtaining a specific, estimated quantitative value previously associated to that range.successful-quantitative-cyber-risk-analysis.

15 Lenny Zeltser, "How CISOs Can Justify Cybersecurity Purchases," Help Net Security, Feb 4, 2020 - https://www.helpnetsecurity.com/2020/02/04/justify-cybersecurity-purchases/.

$$Return_{Control} = Measured\ Effectiveness\ of\ Control * Monetary\ Cost\ of\ Control$$

$$Return_{Validation\ Program} = \frac{\sum_{n=1}^{N} New\ Return_{Control_n} - Total\ Validation\ Costs}{\sum_{n=1}^{N} Inital\ Return_{Control_n}}$$

Figure 3 – Formulas for control and validation program returns

yield, nominal returns, loss reduction and loss avoidance[16], and return on investment (ROI).[17] It must first be acknowledged that security controls are a cost, and so in the truest sense, there is no return from the sole act of improving the efficacy of a control. A control that is not operating efficiently is in essence operating at a loss, and improving efficacy is in reality a loss reduction or a loss avoidance technique as there is no true financial return given the control is still an expense. These terms are commonly if inaccurately interchanged, so be cognizant of how you use them. That noted, these business measures can still unlock the ability to very simply represent each security technology's current monetary value to the organization.

For example, using the first formula in figure 3, if a technology cost US$200K and is only 35 percent effective against tests run during an initial assessment, then one could postulate that the organization is only achieving a US$70K return from that control, a loss of US$130K. If after the control configuration is improved and another assessment is conducted and then the technology's effectiveness is improved, then clearly the control offers more value. In this case, if the effectiveness improves to 85 percent, then the return for that control is US$170K and the loss is reduced to US$30K. In this simple case, and disregarding the cost of the validation activity, this could also be represented as a US$100K and an over 242 percent improvement in effectiveness.

Demonstrating the value of a validation program can be similarly calculated from its overall efficacy improvements. Using the second formula in figure 3, one can sum these improved returns over any number of controls, subtract the total cost in technology and time for validation testing, and divide by the sum of the initial returns of all controls. So, including the previous example, if a second control cost US$500K, had an initial measured effectiveness of 45 percent, a moderately improved effectiveness of 50 percent, and a total validation cost of US$150K, then the total improvement from the security validation investment across both controls with the validation costs deducted is 91.5 percent.[18]

Optimization doesn't have to occur at once; a strategy of leveraging security validation to iterate improvements over time should provide the means to demonstrate a growing yield in tandem. This progress is easily graphed as trend lines and bar charts, both clear and concise forms for demonstrating to executive leaders the increased efficacy of defenses over time. A programmatic approach like this offers security professionals a solid starting point for demonstrating they are good stewards of security budgets for their organizations.

Security validation technologies provide the means to measure the true effectiveness of mitigations against real attacks, providing quantitative measures that can be used for demonstrating the value of investments and a realistic portrayal of the returns from a validation program's ability to bolster the security of the organization. Explaining improved effectiveness using these approaches can help establish the leader as being business oriented, optimizing their pre-existing security technology investments to the maximum extent possible to stop realistic threats applicable to their organization while investing wisely in new solutions to fill discrete gaps in their defenses.

Conclusion

Next generation security leaders "need to be something different: an influential voice in business strategy, technology decisions, and enterprise risk management."[19] These leaders will develop security programs that demonstrate strong effectiveness of their security investments against real threats. They will implement a strategy of focusing proactively on threats most relevant to their organization, continuously tuning their controls to maximize their efficacy, training like their adversaries attack them, and using quantitative measures to justify their investments and demonstrate sound business acumen to their executives and boards. Ultimately, "keeping an eye on the changing risk landscape allows an organization to focus on mitigating [its] most important and relevant risks, while reducing time and resources spent on less important and relevant issues."[20] To that end, security validation technologies will play a key part in further evolving security from a technical field to a more business-oriented discipline in the future.

About the Author

Matt Hartley, CISSP, is a senior vice president at FireEye, Inc., the intelligence-led security company. Matt has spent over 20 years innovating in cybersecurity, threat intelligence, cyber warfare, and information operations in support of corporations and governments worldwide countering the most advanced cyber threats. He may be reached at matt.hartley@fireeye.com.

16 See Will Kenton, "Risk Control," Investopedia, Aug 12, 2019 - https://www.investopedia.com/terms/r/risk-control.asp.

17 See Isaac Kohen, "How to Calculate Your Return on Security Investments," CSO, Oct 2, 2017 - https://www.csoonline.com/article/3229887/how-to-calculate-your-return-on-security-investments.html.

18 In order to more quickly demonstrate the value of validation, large security programs may want to consider deducting validation costs from the improved returns of the first few controls until that cost is completely covered, rather than wait to spread the costs across a wide number of controls if it will take a long time to improve them all.

19 Matthew Doan, "Companies Need to Rethink What Cybersecurity Leadership Is," Harvard Business Review, Nov 27, 2019 - https://hbr.org/2019/11/companies-need-to-rethink-what-cybersecurity-leadership-is#.

20 Joshua Goldfarb, "Keeping a Strong Security Metrics Framework Strong," Dark Reading, Feb 11, 2020 - https://www.darkreading.com/threat-intelligence/keeping-a-strong-security-metrics-framework-strong-/a/d-id/1336962.