

A Data-driven Approach to Rating Cybersecurity Risk and Investing SOC Resources Efficiently

Anup K Ghosh
akghos@gmail.com

Abstract—One of the hardest challenges for companies and their officers is determining how much to spend on cybersecurity and the appropriate allocation of those resources. Security “investments” are a cost on the ledger, and as such, companies do not want to spend more on security than they have to. The question most boards have is “how much security is enough?” and “how good is our security program?” Most CISOs and SOC teams have a hard time answering these questions for a lack of data and framework to measure risk and compare with other similar sized companies. This paper presents a data-driven practical approach to assessing and scoring cybersecurity risk that can be used to allocate resources efficiently and mitigate cybersecurity risk in areas that need it the most. We combine both static and dynamic measures of risk to give a composite score more indicative of cybersecurity risk over static measures alone.

I. INTRODUCTION AND BACKGROUND

Defining and executing an effective security program remains a challenge for all but the most sophisticated enterprise-class Security Operations Centers (SOCs). The naive approach is often buying a set of tools and then expecting the tool stack will solve cybersecurity problems. While tools are necessary for an effective cybersecurity program, they are not sufficient. Like other areas of business, cybersecurity programs need to define the goals of the program, measure the current state and progress, and ensure that the broad areas of business risk are addressed.

Frameworks such as the NIST Cybersecurity Framework (CSF) provide a useful starting point for establishing a cybersecurity program [1]. The NIST CSF establishes five core areas of cybersecurity: Identify, Protect, Detect, Respond, and Recover. Within each of these areas, an organization can identify processes and tools with categories and sub-categories. As an example, Asset Management would be a category within Identify, Access Control within Protect, Continuous Monitoring within Detect, Mitigation within Respond, and Recovery Planning within Recover. Frameworks such as NIST CSF provide a systematic approach to building a cybersecurity program that will also identify gaps in coverage and, therefore, potential risk areas. Following the NIST CSF provides a framework under which organizations can mature their security program

with increasing capabilities. NIST CSF will not necessarily reveal how effective an organization’s cybersecurity program is, or what residual risk lies after following the framework; rather following NIST CSF will likely ensure the foundational building blocks of a cybersecurity program are in place and provide a roadmap for maturing a program over time.

Taking a different approach is the MITRE ATT&CK framework for understanding adversarial tactics, techniques and protocols (TTPs) [2]. MITRE ATT&CK takes a threat-informed defense approach to understand the types of TTPs adversaries use against targets and builds a security program accordingly. SOC teams can use MITRE ATT&CK as a reference for ensuring coverage of attacks they are likely to experience, given the toolsets they have. For example, If APT28 is a threat actor likely to target you, the MITRE ATT&CK matrix will show the current TTPs APT28 employs. By studying the coverage of tools in your environment against those TTPs, an organization can determine if they have adequate tool coverage against their attack types. Organizations generally need a level of sophistication to gather threat intelligence and also to understand their coverages from their tool stacks in order to effectively use MITRE ATT&CK. By itself, MITRE ATT&CK only provides guidance, but will not be able to determine how effective your security program is, or where your residual risk is after you have employed an ATT&CK based defense.

The author previously developed a SOC Capability Progression model that is intended to bridge the gap between foundational security programs such as the NIST CSF and advanced security capabilities like the MITRE ATT&CK framework [3].

As shown in Figure 1, the SOC capability progression model, organizations can stepwise level up their capabilities starting with foundational capabilities (L1) to establishing security investigation infrastructure (L2), security automation (L3), advanced analytics (L4), to predictive analytics (L5) for advanced organizations. While this capability progression model lays out a roadmap for building capabilities, like the other frameworks, it does not provide a means to measure how effective a security program is in mitigating business risk due to cybersecurity concerns.

While these are not risk-driven models, they are the prevailing models on which security operations programs are built today. They are useful for defining a security program and understanding what gaps need to be addressed. However, even adherence to one model of another does not provide a measure of how well one has secured a network against attack nor




Fig. 1. SOC Capability Progression Model (see <https://www.accenture.com/us-en/blogs/security/level-up-soc-game-one-logical-step-at-a-time>)

what residual risk is present. One of the challenges in security operations is while several frameworks including the ones mentioned above have been developed to progressively build more mature security programs, there is poor linkage to how secure or risky an organization’s network is even after they have built to a certain level of maturity in a given framework.

Separately an academic body of work exists that have built economic and mathematical models of cyber risk and optimal investment. The Gordon-Loeb Model is one such example that builds up a mathematical approach for optimal cyber investment [4]. The model is parameterized by variables such as the probability of breach and the expected loss if a breach were to occur. In practice, most organizations are challenged to even know where their assets are, let alone the expected cost of loss of data, or their probability of their breach. These models while appearing mathematically sound are difficult to put into practice in a security program run by security practitioners. The opportunity is to marry a practical security risk scoring with a security program that measures security as an organization improves its maturity.

II. SCORING RISK

Recently, companies are building scoring approaches to security risk, including Security Scorecard and Microsoft. Although there is no canonical scoring algorithm or approach yet, we believe the beginnings of a widely accepted score may involve different approaches to assessing risk until actuarial science proves it is sound.

The basic scoring methodology of these approaches involves determining a set of risk factors, their relative importance on risk via a weighting function, and an algorithm for calculating a weighted scoring average. The scoring methodology for each

factor will differ depending on which factors are considered and the relative importance assigned to each factor.

Security teams will differ over what factors are relevant and the relative importance of each factor. For instance, many of the factors from Security Scorecard involve security best practices or hygiene, such as a malformed SPF record, self-signed or an expired certificate and weak ciphers on TLS protocols. Others may involve exposed insecure services on the Internet or vulnerable unpatched services. While one can argue the relative importance of each factor, from a SOC operations point of view, the important idea is to identify the priority areas of risk exposure and then act accordingly to remedy them. Scoring is both a gamification of remediating risk as well as a means to provide relative comparisons to other similar firms. This provides one answer to “how much is good enough?” from a business point of view. From a scoring perspective, the weightings should be adjustable by security experts to meet the context of the business and its risk factors.

Today most commercial scoring methodologies are both static and external. With static scanning, a moment-in-time scan of assets is taken to populate the scoring fields. External scans are conducted using Internet-based scans of a company’s visible infrastructure, and often without their permission. While an outside scan is useful, a risk scoring methodology that captures only the external attack surface at a moment in time is limited at best and a poor approximation of risk at worst. The frequency of the scans determines the window between when there is an exposure and when the company is informed about it. As an example, if a scan is conducted on a monthly basis, and exposure due to a vulnerability is released immediately afterward, then the exposure window may be 30 days after the scan is performed, barring other

means of vulnerability discovery in the interim. External scans can be useful in identifying vulnerabilities and open doors, but today the commercial external scanners are conveying mostly security best practices (or the lack thereof) and security hygiene. As a result, they are often dismissed for their lack of usefulness by security teams.

Internal scans, on the other hand, can be far more useful at identifying vulnerabilities that an attacker may exploit once they obtain a beachhead on a network. For most adversaries, getting a point of presence on a network, i.e., the initial compromise of a machine, is as easy as launching a phishing campaign against a target company. As soon as an adversary gains access to the network, they will be able to scan it from the inside as effectively as they can from the outside to find targets of interest. It would be negligent for SOC teams not to scan their own networks and not take into account the risk associated with vulnerable internal assets including servers, desktops, and other networked devices including network, security, and Internet of Things devices.

Dynamic approaches measure the observed behavior of the network and naturally complement the static scans. Network and endpoint monitoring is an example of dynamic measurement of behavior. Observing suspicious activity on the network or, conversely, the absence of suspicious activity informs how effective an estimate of risk the static scans are. Internal monitoring can also reveal suspicious behavior from insider threats. For example, detecting large uploads of files to an IP address not normally used by a business is an indicator of insider threat and data leakage. Given effective observation mechanisms, the lack of suspicious activity on a network may belie a high-risk score or the presence of numerous suspicious activities belies a low-risk score. Thus, a risk score as a composite of both static and dynamic measurement is a more effective estimate of true risk on the network.

III. RISK CATEGORIES & SCORING

Security operations teams will often spend most of their focus and resources on detecting and thwarting intruder threats on the network. While this is a key focus, businesses face additional risks that fall within the remit of security operations teams. Ignoring these risks comes at the peril of the business. Aligning security to the business needs requires concerted management of these risks. We identify key security risks that businesses need to manage in Table 1.

The categories in Table 1 are not complete and will likely grow especially as new risk evolves. For instance, it is now understood that supply chain risk is a key risk for businesses to manage after some notable supply chain compromises over the last two years. One might also include Operational Technology (OT) as a risk category for companies that have manufacturing OT networks. The key focus of detecting threats on the network is addressed in the security monitoring category, but a SOC team would be remiss if they were not paying attention to the security of the company’s cloud assets (cloud security) and of bespoke software applications (app sec) they publish. Likewise an inventory of devices owned

TABLE I
SECURITY RISK CATEGORIES

Categories of Risk	Description
App Security	Vulnerability discovery in bespoke applications on network usually requiring source code analysis
Asset Discovery/Identification	Identify all devices with IP addresses on network, infer function based on protocol analysis
Cloud Security	Vulnerability discovery and security misconfiguration of cloud assets
Compliance	Conformance to PCI, PHI, PII, SOC compliance standards
IT Policy	Conformance to company IT policy for endpoints and servers
SaaS Security	Security configuration of approved and identification of unapproved SaaS services
Supply chain/SBOM	Compromise of key software libraries or programs from suppliers and open source
Suspicious Activity	Monitoring of networks and endpoints for suspicious activity
Threat Intelligence	Dark web monitoring for compromised credentials, hacker chatter, IP reputation, and breach notification of vendors/suppliers
Vulnerability Management	Identification of software vulnerabilities in software, internal, externally facing and cloud services

by the business (asset discovery and identification) as well as their vulnerability posture (vulnerability management) is foundational to understanding risk in the enterprise. With so much of IT being outsourced to SaaS providers (SaaS Security), the configuration and security of third party services that contain business critical data is an increasingly important exposure area. Finally, threat intelligence is part and parcel to every serious security program. Understanding when accounts are compromised, when a company is targeted, when vendor suppliers are compromised, and whether your company is being targeted is a key risk consideration.

Scoring risk in each of these categories remains a complex challenge, one we will not address in this paper. Rather, we present the framework from which to rate risk and allow the field to develop increasingly sophisticated risk models, of which a lot of academic approaches exist. The goal for security practitioners is to make risk understandable and actionable. Risks should tie to addressable issues from which actions can be taken to reduce risk. Our criteria for a good rating scheme is an understandable scale, consistency across different firms and networks, and linearity of ratings as a function of risk. In other words, when an issue is identified, a rating computed, and the risk redressed, the change in rating should have a proportional impact to the issue addressed.

Our approach is to use a risk/resiliency rating of 0 to 100 in each category. A perfect score of 100 indicates no observed indicators of risk. Each identified risk detracts from 100 proportional to the severity of the risk. An overall score




Fig. 2. Risk Scores Computed Over Time

is computed across all risk categories on a scale of 0 to 100 by using a weighted average function. If all categories were equally weighted it would simply be the average of all categories. The weightings themselves can be changed according to the industry segment, size, and threat profile. For example, a tech or SaaS firm may more heavily weigh cloud security, while a regulated industry may more heavily weigh compliance.

The computation in each specific risk category will necessarily vary from category to category. For example, when assessing vulnerabilities, we use the product of the vulnerability’s severity rating, e.g., its CVSS score, its exploitability index, e.g., its EPSS score, and its discovered frequency in the enterprise to bubble up the highest priority vulnerabilities to address to minimize risk.

While risk scores are typically computed from scan data, each scan represents a sampling of the security state. Collected over time, they represent a time series for risk from which risk estimates as a function of time can be plotted. Figure 2 shows computed risk from a live production network as a time series. By examining the scan at any point in time, one can assess why a score moves up or down at a moment in time. A vulnerability discovered through a scan in Figure 2 dropped the rating from a 75 to a 66 until it was remediated. The framework provides flexibility for the best algorithms to compute risk to emerge from different researchers. Experience will reveal what the best algorithms are over time based on computed risk versus actual risk.

IV. CONCLUSIONS AND FUTURE DIRECTIONS

Security has become a primary business risk for businesses of all sizes and in all industries. As such, security professionals and security operations teams need to co-opt business risk based approaches to align security with business objectives. Most areas of a business today are managed by numbers

and project managed with Key Performance Indices (KPIs). Increasingly executive management and boards will require management by measure from security leaders and by extension security teams. A quantitative approach to understanding and acting on cybersecurity risk across the enterprise will begin to place security as an understood business process on the same level as other business areas such as sales, marketing, software development, and manufacturing.

In this paper, we outline a practical approach for identifying and rating risk across the enterprise that can be put into practice by security teams. While the sheer breadth of the categories and the depth of factors within each category may seem onerous, our experiences show that data can be collected and computed through automation for each factor in each risk category. In other words, the collection and computation can be automated through software. The clear benefit of computing a score is security leaders will be able to identify areas that need improvement and more resourcing. A quantitative approach allows comparison with other similar sized firms and in similar segments. It also supports gamifying the security work so that individual contributors can be recognized for their contributions to security. By capturing risk scores as a time series, one can plot trending and show improvement – or not – against baseline scores. Management can create measures of effectiveness for teams and individual contributors, Scores computed across companies will facilitate statistical measures of security for industry segments. In turn, this will support computing deciles of performance for individual companies much like quality metrics have been applied to manufacturing programs for decades.

In terms of future directions, we will publish our results from applying this approach to security programs in different size companies to put to real-world test the approach and scoring. We continue to work on the specific algorithms to

scoring in the individual risk categories and we expect these will evolve over time. Finally, we expect that as this approach becomes more widely adopted by both industry and standards this will facilitate maturing security programs of different sized companies across multiple sectors.

ACKNOWLEDGMENT

The author would like to thank Stan Ivanov for his contributions to these ideas and the framework in addition to Figure 2 which was computed from live data on a production network. The author thanks Allison Lewis for her editing and revisions to the manuscript.

REFERENCES

- [1] NIST Cybersecurity Framework. <https://www.nist.gov/cyberframework>
- [2] MITRE ATT&CK Framework. <https://attack.mitre.org/>
- [3] SOC Capability Progression Model. <https://www.accenture.com/us-en/blogs/security/level-up-soc-game-one-logical-step-at-a-time>
- [4] Gordon, L.A. and Loeb, M.P. (2002) The Economics of Information Security Investment. *ACM Transactions on Information and System Security*, 5, 438-457.