ASEE AMERICAN SOCIETY FOR ENGINEERING EDUCATION

Paper ID #45008

Threat Modeling for Optimal Enterprise Protections Against Known Cybersecurity Threats

Mr. Branko S. Bokan, The George Washington University

Branko Bokan is a PhD candidate at the School of Engineering and Applied Science, George Washington University under professor Joost Santos.

Branko is a Cybersecurity expert at the Cybersecurity and Infrastructure Security Agency (CISA) at the Department of Homeland Security (DHS). In his professional role he is responsible for defending the Federal Civilian Executive Branch of the U.S. government against cyber threats and building a cyber resilient federal enterprise.

Dr. Joost R. Santos, The George Washington University

Threat Modeling for Optimal Enterprise Protections Against Known Cybersecurity Threats

Abstract— To prioritize limited resources available to protect against cybersecurity attacks, organizations must adhere to risk management practices. These in turn necessitate a proper framing of risk, which requires a 'set of triplets' to be understood – (i) a scenario in which a threat exploits a vulnerability, (ii) the likelihood, and (iii) the impact of the scenario taking place. While the other elements of risk triplets are relatively easy to assess, the threat factor remains the most elusive. The traditional threat modeling methodologies work well on a small scale — but do not allow for a comprehensive evaluation of an entire enterprise architecture to identify gaps where cybersecurity protections do not exist and where future investments are needed. They also do not enumerate and consider all threats observed in the wild. A new threat modeling approach – the Cybersecurity Architecture Review – allows decision-makers to select cybersecurity capability portfolios that maximize protections against known cybersecurity threats. It allows organizations to look at their cybersecurity protections from the standpoint of an adversary and allows them to evaluate entire cybersecurity architectures. The methodology scores cybersecurity capabilities for their ability to protect against all threats that were previously observed in the wild and enumerated using a cyber threat framework such as Mitre ATT&CK. The results form a matrix called capability coverage map – a visual representation of protections coverage, gaps, and overlaps against threats. This paper provides a proof of concept for proposed future research to determine how commonly applied cybersecurity capabilities protect against known threats, whether organizations use the most efficient portfolios of capabilities, whether these portfolios are selected based on the actual threat landscape or vendor pressure, how different demographics perceive their protection coverage, and to what extent those common protections overlap.

Keywords: threat modeling, cybersecurity, cybersecurity architecture, cybersecurity capabilities, cyber threat framework, risk, cybersecurity risk, risk management.

I Managing cybersecurity risk

Cybersecurity risk is described as a function of an adverse event (or a scenario), impact, and likelihood of that event. [1] To fully understand risk, one needs to consider both the threats and vulnerabilities concurrently as for an adverse scenario to occur, a threat needs to exploit a vulnerability. Previously, there was no well-documented and accepted methodology that allowed for proper consideration of the threat component when making risk-based decisions on selection and deployment of cybersecurity capabilities. The traditional threat modeling methodologies were hard to apply to an entire infrastructure at an enterprise level, leaving organizations with

consideration of only the vulnerability component of the scenario when managing cybersecurity security risk.

The inability to incorporate the threat component to describe a potential adverse event and ultimately inform risk management decisions resulted in inadequate infrastructure protections, capability "blind spots," wasted limited resources on capabilities that do not cover actual threats organizations are facing, or multiple capabilities covering the same limited threats.

An emerging threat-informed approach – Cybersecurity Architecture Review – allows organizations to optimize the capability coverage against known cybersecurity threats that have been observed in the wild. In the context of this article, capability coverage is described as a measure of the ability to protect against, detect, or respond to a threat action and it can be expressed as limited, moderate, significant, not applicable, and none. The optimization is achieved by determining the coverage of the existing capabilities, identifying gaps and needs for future capabilities, and flagging overlaps – areas where multiple capabilities protect against the same types of threats thus unnecessarily multiplying the cost.

To assess the risk associated with a certain scenario, Kaplan and Garrick suggest answering the following three questions: a) what can go wrong; b) what is the likelihood; and c) if it does happen, what are the consequences? [2] The risk assessment allows the decision-makers to enumerate and prioritize risks in order to make informed risk-based (response) decisions. During the assessment, a researcher identifies specific threat sources and threat events the sources could produce, identifies vulnerabilities that could be exploited by those sources, determines the likelihood of exploitation, and determines potential adverse impact. In the final step, the researcher quantifies the risk as a function of the likelihood of vulnerability being exploited by the threat and the impact of the exploitation. [3]

Risk managers frequently attempt to measure and prioritize risk by using risk matrices to map the frequency of occurrence and some arbitrary severity rating and express risk as a number or a qualitative value such as "low, medium, moderate." This practice may also be encouraged by mid- and top-tier managers who may demand the risk be described in simple terms or as a single number. Cox [4] warns that such practices may prove to be "worse than useless" because they provide poor resolution (can only correctly compare a small fraction of pairs of hazards), are erroneous (can mistakenly assign different ratings to the same risks), provide suboptimal resource allocation (matrices do not provide effective countermeasures), and have ambiguous inputs and outputs (different researchers may interpret the same inputs and outputs differently). Kaplan [2] also argues that thinking of risk as a single number as a product of two values (probability times consequence) is very misleading and points to the same issues as Cox. Instead, Kaplan suggests that one should think of risk as a "family of curves" since not a single number, not a single curve, but only a series of curves can properly communicate all possible sets of

triplets. While there is no single good way to measure, prioritize, and express risk – risk matrices remain the most popular and widely adopted approach and therefore one needs to exercise caution when using them. [4]

Haimes [5] introduces three additional questions to help guide the risk response:

a) What can be done?
b) What options are available and what are their associated trade-offs in terms of all costs, benefits, and risks?
c) What are the impacts of current management decisions on future options?

Determining what can be done and what options are available generally falls into one of five major risk response strategies: accept, avoid, mitigate, share, and transfer risk. [6] Mitigation is associated with activities that result in reduction of either the likelihood or the impact of the adverse event. In cybersecurity, these activities may include the implementation of countermeasures or cybersecurity controls, deployment of cybersecurity capabilities, policies, or reduction of the attack vector (e.g., by applying updates to vulnerable software).

For an adverse event or a scenario to exist, a threat needs to exploit a vulnerability. A threat can be “any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, or the Nation through an information system via unauthorized access, destruction, disclosure, modification of information, and/or denial of service.” [7] A vulnerability, in turn, is a “known weakness in a system, system security procedures, internal controls, or implementation by which a [threat] actor or event may intentionally exploit or accidentally trigger the weakness to access, modify, or disrupt normal operations of a system-resulting in a security incident or a violation of the system's security policy.” [7]

II Threat modeling

Threat modeling is a risk assessment approach that informs the risk management process and allows for enumeration, analysis, and prioritization of threats to and vulnerabilities in an information system. The results of threat modeling inform decisions on which threats and vulnerabilities are associated with the highest risk and need to be mitigated first. Threat modeling can be conducted from the perspective of an asset (something we are trying to protect – e.g., an information system) and from the perspective of an attacker (thinking like an adversary). [8]

One of the first formal threat analyses for information systems was developed by AT&T for the Department of Defense’s Security Vulnerability Analysis (SVA) for System Security Engineering (SSE) process in the early 1980s. [9] The model introduced "threat logic trees" for threat decomposition. While the model set the stage for future efforts in threat modeling, the

biggest downside was that it was dependent on the identification of "potential threats", which was poorly defined and left at the discretion of the researcher.

Significant improvements to threat modeling were made by a group of authors led by Bruce Schneier in "Methodology for Characterizing Attacks and Choosing Rational Countermeasures", a research sponsored by the National Security Agency. [10] The methodology introduced a way to enumerate and visually represent possible attacks (threat actions) and weigh them based on the risk, access, and cost to the adversary through five broad steps. The analysis begins with a root node of a tree representing either the component of the analysis or the objective of the adversary. Child nodes are formed by decomposing the root node first into its lifecycle phases, then each phase into physical security and trust model access categories. Each new node can be further decomposed in the same way until a series of vulnerability leaves is reached. Each leaf is assigned qualitative values for risk, access, and cost to the adversary. Only leaves that need countermeasures (e.g., meet adversary's objectives, capabilities, or offer sufficient return) are preserved. Prioritized mitigations for the exploitable vulnerabilities are determined and deployed. Just like with SVA for SSE described in the previous section, the model leaves it up to the researcher to determine possible threats, which produces greatly different results based on the researchers' backgrounds and collective experience. It does not standardize the enumeration of common threats, nor provide any guidance for their identification.

In 1999, Microsoft Corporation introduced STRIDE – a threat modeling methodology named after major categories of threats occurring in the wild – Spoofing of user identity, Tampering with data, Repudiability, Information disclosure, Denial of Service, and Elevation of privilege. [11] Enumerating all possible threats and vulnerabilities in a single software product, let alone an entire information system or organization, can be a daunting process. The STRIDE methodology simplified this task and assisted security engineers by grouping known threats into six categories [11] and describing various products and services each category applies to. The methodology is applied by decomposing a system under consideration into components, analyzing each component for susceptibility to threats in each category to discover associated vulnerabilities and develop appropriate threat mitigation measures. [12]

The DESIST methodology was developed by Gunnar Peterson and followed a similar process as STRIDE. It was named for Dispute, Elevation of privileges, Spoofing, Information disclosure, Service denial, and Tampering [8]. This model offered a first attempt to group threats into common categories and provide some guidance for consistent analysis across different systems and their components.

Around the same time, another popular threat modeling methodology emerged – the Process for Attack Simulation and Threat Analysis (PASTA). It focuses on business objectives as drivers for both information system requirements and associated security responses. PASTA assumes that

organizations in different industries face different types of threats and therefore only those impacting the organization should be mitigated. [13] The process comprises seven stages with a goal of identifying and addressing the most viable threats to a specific application. [14]

All the traditional threat modeling methodologies discussed so far work well on a limited scale, when evaluating targets such as single data field, a software application, or a system component – but they do not allow for a comprehensive evaluation of an entire enterprise architecture to identify gaps in cybersecurity protections and identify areas where additional protections (future investments) are needed. They also do not enumerate and consider a comprehensive set of threats that have been observed in the wild.

In 2015, the Department of Defense (DOD) developed a novel approach to threat modeling that allowed DOD to apply a threat informed risk management process and, for the first time, look at the cybersecurity architectural capabilities and their ability to protect against the actual threats from the standpoint of an adversary. Named DOD Cybersecurity Architecture Review (DODCAR), the methodology has been widely used by DOD to evaluate the threat landscape, identify gaps where protections do not exist but are necessary to inform the future investments into new capabilities, or to identify capability overlaps to inform decisions to retire redundant capabilities (e.g., two or more different products serving the same purpose and protecting against the same type of threat). The Department of Homeland Security later adopted and further expanded the methodology under the name .govCAR. [15]

III Cybersecurity Architecture Review

The Cybersecurity Architecture Review relies on the cyber threat framework to enumerate possible threat actions (referred to as tactics, techniques, and protocols, or TTPs) carried out by adversaries in cyber-attacks and incidents that have been observed in the wild. The Mitre ATT&CK cyber threat framework represents enumerated threat procedures in a matrix format. It breaks down the procedures into 14 tactics that represent the “why” of a technique – the adversary’s tactical goals or reasons for performing an action. [16] Each tactic is then broken down into individual techniques that describe how an adversary achieves a tactical goal by performing an action. Finally, each technique can be broken down into sub-tactics (a more specific description of adversarial behavior) and procedures (specific implementations the adversary uses for techniques or sub-techniques). Figure 1 shows the entire Mitre ATT&CK Matrix for Enterprise, with 14 tactics at the top and each cell underneath representing a total of 227 techniques.

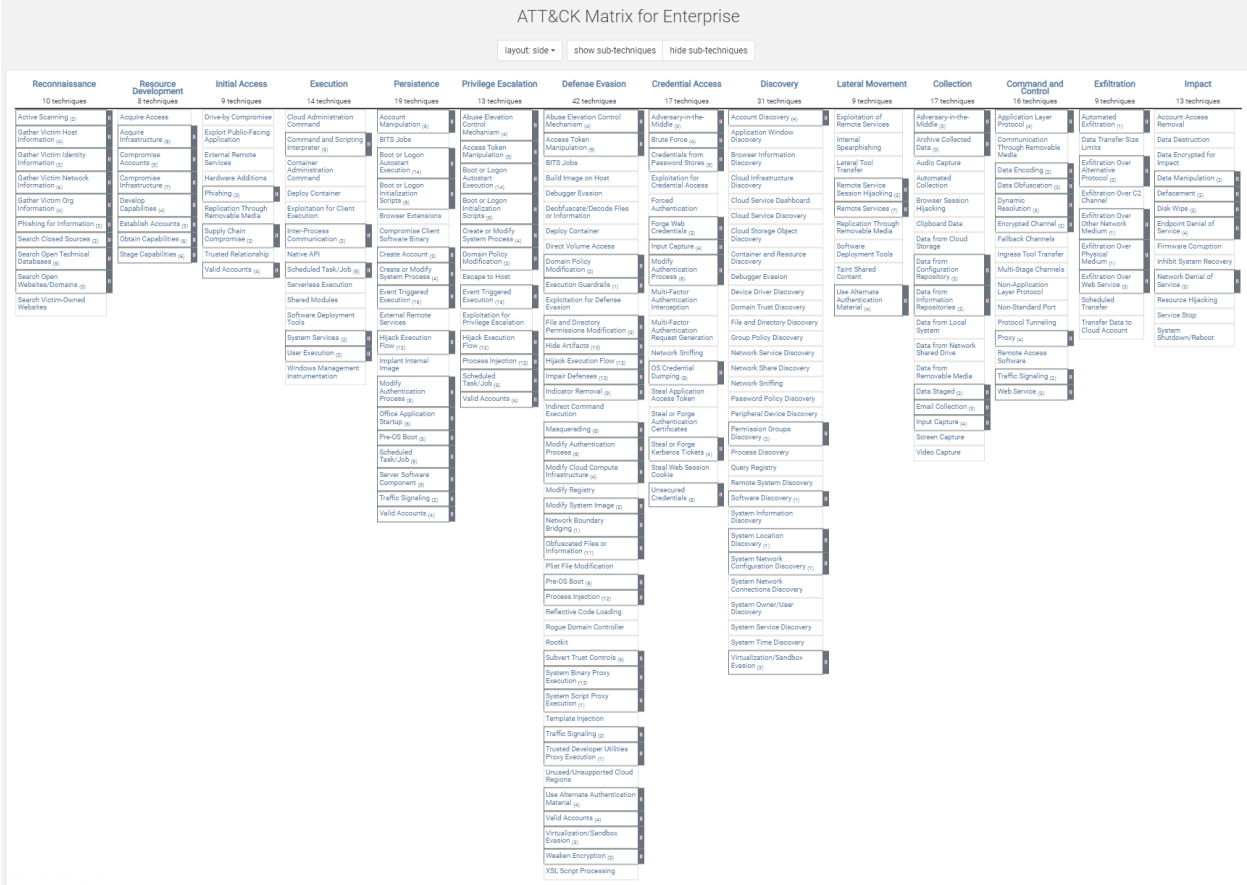

Figure 1 - Mitre ATT&CK Matrix for Enterprise [16]

In the next step, researchers identify and define the attack surface – the target cybersecurity architectures for review. The target architecture comprises computing technologies (computing infrastructure) and defensive mechanisms intended to protect it. Researchers identify capabilities under consideration, their topologies (e.g., positions on the network), and sources of and destinations for network traffic that are routed through those capabilities (network flows). [17] They may also note particular flows of interest that may bypass the capabilities under consideration.

Using an improved approach since the first version [18] of this methodology was published in 2018, the current version allows researchers to decompose capabilities into detailed features – narrow functions that define that capability and constitute its building blocks (e.g., a firewall capability contains a feature that blocks ports and another feature that examines packet content). This approach provides more flexibility and allows for the reuse of scoring data, as individual features can be found in various capabilities (different products by different vendors). [17]

Finally, the enumerated threat actions and defined cybersecurity architectures are arranged into a matrix, with the former listed at the top as column headers and the latter on the left side as the

row titles. The arrangement is similar to the Capabilities Matrix in the Quality Function Deployment (QFD) method [19] that uses matrices to map capabilities to design requirements. However, in the Cyber Architecture Review analysis, the following questions are asked at each intersection of a threat action and corresponding capabilities (and features): a) can this capability (or feature) detect this threat action; b) can this capability protect against this threat action; and c) can this capability help in recovery against this threat action? The questions are aligned with three out of five functions of the NIST cybersecurity framework (CSF – not to be confused with the cyber threat framework). [20] The answers to the questions above, which can be binary (e.g., yes or no) or ranked in some way (e.g., some, moderate, or significant coverage), are entered into the matrix, and the results generate the capability coverage map – a visual representation of capability coverage, gaps, and overlaps against the threats. Colors can be used to represent different correlation strengths (e.g., red for no coverage, green for significant, etc.). Coverage maps for multiple capabilities can then be combined by overlaying them on top of each other to demonstrate the coverage of the entire organizational defense-in-depth architecture. Figure 2 shows a visualization of a coverage map for a sample cybersecurity capability where colors denote coverage from green for significant to red for no coverage, and dark gray for no mapping (not applicable).

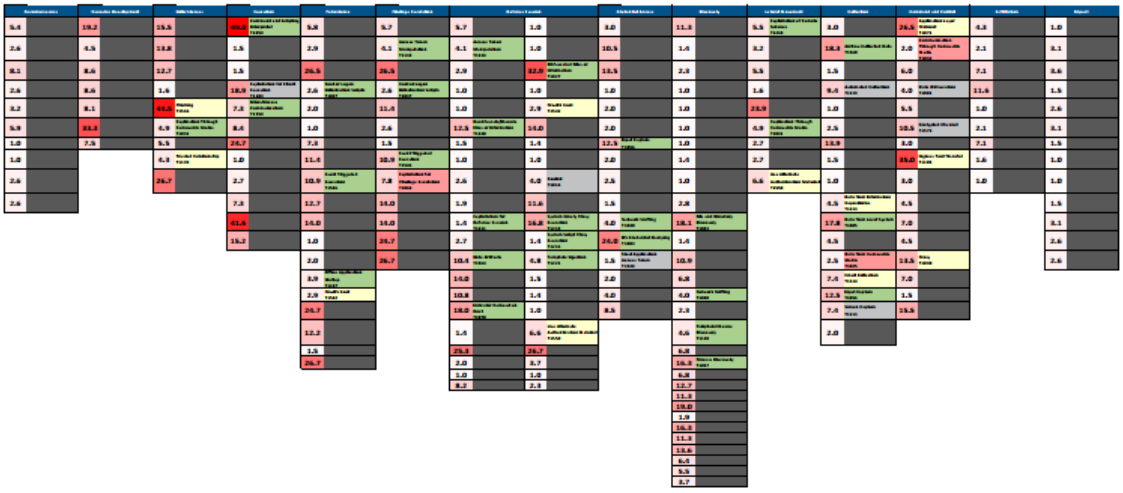

Figure 2 - Sample visualization of capability coverage map [17]

To allow for their prioritization, threat actions can be ranked based on their frequency of occurrence in the wild and maneuverability (the number of different threat actions that can be used to achieve the same objective). The ranking results can be plotted on a threat heat map to visually highlight threat actions based on their priority. The heat map can be overlaid on top of any coverage map to prioritize future capability focus.

IV Proof of concept

Several regulatory and industry compliance mandates impose various types of cybersecurity countermeasures, known as controls, that are required for information systems that process and store data related to a particular industry. The regulators face the same challenges as organizations – no previously documented methodology allowed them to make threat informed risk decisions on the best control portfolios. The Cybersecurity Architecture Review threat modeling approach can be applied to the analysis of coverage of administrative cybersecurity controls (or non-materiel capabilities) in the same way researchers previously applied them to technical (materiel) capabilities.

The Federal Risk and Authorization Management Program (FedRAMP) is a federal program that promotes the adoption of secure cloud services across the federal government by providing a standardized approach to security and risk assessment for cloud technologies and federal agencies. [21] To ensure a standardized and repeatable approach to securing cloud-based offerings used in the federal government, FedRAMP defines baselines for cybersecurity controls that must be applied in cloud-based information systems based on their risk categorization. The deployment, maintenance, and monitoring of a large number of cybersecurity controls require significant resources, and their effectiveness is hard to measure.

In an effort to prioritize the government's cybersecurity investments, to utilize resources effectively, and to reduce the greatest amount of risk, FedRAMP applied the Cybersecurity Architecture Review methodology to analyze the effectiveness of and prioritize the NIST SP 800-53 [22] cybersecurity controls based on their ability to protect against real-world threats. The results of the threat modeling effort combined with the cyber threat framework heat maps in Figure 3, allowed the FedRAMP program to establish an assessment threshold and prioritize the evaluation of only those security controls which fall above the established threshold, eventually providing a foundation for streamlining the security authorization process. [23]

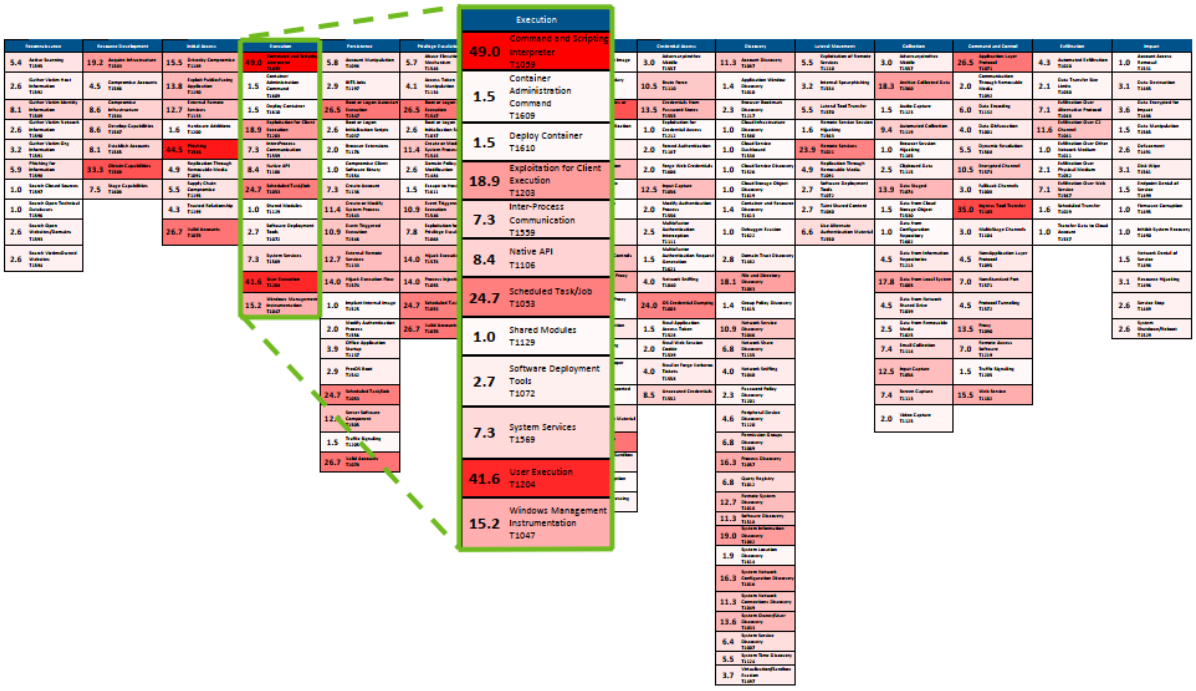

Figure 3 - Cyber Threat Framework heat map [17]

By mapping the threat modeling results to cybersecurity capabilities [24] FedRAMP also improved their methodology for measuring an overall prioritized threat-based risk profile (maturity level) of cloud-based information systems.

This proof of concept demonstrated that by applying the Cybersecurity Architecture Review threat modeling methodology, cloud service providers and their customer organizations can prioritize their resources and focus their efforts on deployment of security controls that mitigate threats and vulnerabilities that pose the most risk to information systems and data, eventually delivering more secure systems that protect against the most current threat landscape. Additionally, the resulting risk profiles allow organizations to strategically manage the protection of their systems.

V Recommendations for future research

The novel approach to threat modeling based on the Cybersecurity Architecture Review methodology presents opportunities for future research to help understand what drives decision-makers to select the current cybersecurity capability portfolios (both materiel such as cybersecurity technologies and non-materiel such as cybersecurity controls and policies), how existing portfolios protect organizations against actual cybersecurity threats observed in the wild,

and how to maximize the coverage of portfolios in terms of their ability to protect against, detect, and respond to cybersecurity attacks.

REFERENCES

[1] National Institute of Standards and Technology, "Federal Information Processing Standards Publication (FIPS 200): Minimum Security Requirements for Federal Information and Information Systems," Gaithersburg, 2006.

[2] S. Kaplan and B. J. Garrick, "On the Quantitative Definition of Risk," *Society for Risk Analysis,* pp. 11-27, 1981.

[3] National Institute of Standards and Technology, "Special Publication 800-30: Guide for Conducting Risk Assessments Revision 1," Gaithersburg, 2012.

[4] L. A. Cox Jr, "What's Wrong with Risk Matrices?," *Risk Analysis,* pp. 497-512, 2008.

[5] Y. Y. Haimes, "Total Risk Management," *Risk Analysis,* pp. 169-171, 1991.

[6] National Institute of Standards and Technology, "Risk Management Framework for Information Systems and Organizations - A system Life Cycle Approach for Security and Privacy NIST SP 800-37 Revision 2," Gaithersburg, 2018.

[7] Committee on National Security Systems Glossary, "CNSSI-4009 Committee on National Security Systems (CNSS) Glossary," Ft. Meade, 2022.

[8] A. Shostack, Threat Modeling: Designing for Security, Germany: Wiley, 2014.

[9] J. D. Weiss, "A System Security Engineering Process," in *14th National Computer Security Conference - Information Systems Security: Requirements and Practices* , Washington, DC, 1991.

[10] B. Schneier, C. Salter, S. Saydjari and J. Wallner, "Toward a secure system engineering methodology," in *7th New Security Paradigms Workshop Proceedings*, CHARLOTTESVILLE, VA, 1999.

[11] L. Kohnfelder and P. Garg , "The threats to our products," April 1999. [Online]. Available: https://cloudblogs.microsoft.com/microsoftsecure/2009/08/27/the-threats-to-our-products/.

[12] S. Hernan, S. Lambert, T. Ostwald and A. Shostack, "Uncover Security Design Flaws Using The STRIDE Approach," *MSDN Magazine - The Microsoft Journal for Developers,* 2006.

[13] Versprite, "PASTA Threat Modeling," 1 December 2020. [Online]. Available: https://versprite.com/tag/pasta-threat-modeling/.

[14] T. UcedaVelez and M. M. Morana, Risk Centric Threat Modeling - Process for Attack Simulation and Threat Analysis, Wiley, 2015.

[15] B. Bokan and J. Santos, "Threat Modeling for Enterprise Cybersecurity Architecture," Washington, 2022.

[16] The MITRE Corporation, "Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) Matrix for Enterprise," 1 November 2023. [Online]. Available: https://attack.mitre.org/.

[17] The Department of Homeland Security, ".govCAR Methodology Version 3.0," Washington, 2022.

[18] The Department of Homeland Security, ".gov Cybersecurity Architecture Review (.govCAR) Methodology," Washington, 2018.

[19] J. Wasek, S. Sarkani and T. Mazzuchi, "Measuring Defense Acquisition Capabilities with QFD," *Military Operations Research,* p. 75–92, 2009.

[20] National Institute of Standards and Technology, "The NIST Cybersecurity Framework 2.0 - Initial Draft," Gaithersburg, 2023.

[21] General Services Administration (GSA), "fedramp.gov," 25 March 2024. [Online]. Available: fedramp.gov.

[22] National Institute of Standards and Technology, "NIST Special Publication 800-53 Rev. 5 Security and Privacy Controls for Information Systems and Organizations," Gaithersburg, 2020.

[23] General Services Administration (GSA), Federal Risk and Authorization Management Program (FedRAMP), "Threat-Based Risk Profiling Methodology," Washington, 2022.

[24] National Institute of Standards and Technology, "Interagency Report (NISTIR) 8011 Automation Support for Security Control Assessments, Volume 1," Gaithersburg, 2017.

[25] N. Shevchenko, "Threat Modeling: 12 Available Methods," 3 December 2018. [Online]. Available: https://insights.sei.cmu.edu/sei_blog/2018/12/threat-modeling-12-available-methods.html.

[26] Lockheed Martin Corporation, "The Cyber Kill Chain," 1 November 2023. [Online]. Available: https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html.

[27] National Security Agency, "NSA/CSS Technical Cyber Threat Framework v2," National Security Agency, Washington, 2018.