Threat Modeling for Enterprise Cybersecurity Architecture

Branko Bokan (brankobokan@gwu.edu) and Joost Santos (joost@gwu.edu),
The George Washington University, Washington, DC, USA

2022 Systems and Information Engineering Design Symposium (SIEDS) | 978-1-6654-5111-6/22/$31.00 ©2022 IEEE | DOI: 10.1109/SIEDS55548.2022.9799322

***Abstract*— The traditional threat modeling methodologies work well on a small scale, when evaluating targets such as a data field, a software application, or a system component—but they do not allow for comprehensive evaluation of an entire enterprise architecture. They also do not enumerate and consider a comprehensive set of actual threat actions observed in the wild. Because of the lack of adequate threat modeling methodologies for determining cybersecurity protection needs on an enterprise scale, cybersecurity executives and decision makers have traditionally relied upon marketing pressure as the main input into decision making for investments in cybersecurity capabilities (tools). A new methodology, originally developed by the Department of Defense then further expanded by the Department of Homeland Security, for the first time allows for a threat-based, end-to-end evaluation of cybersecurity architectures and determination of gaps or areas in need of future investments. Although in the public domain, this methodology has not been used outside of the federal government. This paper examines the new threat modeling approach that allows organizations to look at their cybersecurity protections from the standpoint of an adversary. The methodology enumerates threat actions that have been observed in the wild using a cyber threat framework and scores cybersecurity architectural capabilities for their ability to protect, detect, and recover from each threat action. The results of the analysis form a matrix called capability coverage map that visually represents the coverage, gaps, and overlaps against threat actions. The threat actions can be further prioritized using a threat heat map – a visual representation of the prevalence and maneuverability of threat actions that can be overlaid on top of a coverage map.**

The paper discusses the new threat modeling methodology and proposes future research with a goal to establish a decision-making framework for selecting cybersecurity architectural capability portfolios that maximize protections against known cybersecurity threats.

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

I. Introduction

Threat modeling is a structured process for enumeration, analysis, and prioritization of threats to, and vulnerabilities in, an information system [1]. The results of such a process can be used to inform decisions on which threats and vulnerabilities are associated with the highest risk, and which cybersecurity capabilities are required to address them. The process can be accomplished from two different perspectives: from the perspective of an asset or a system (something we are trying to protect) and from the perspective of an attacker (thinking like the adversary) [1].

This process works well on a small scale—it can be easily applied to a single data field, a software application, or a system component but it does not scale well and fails when we try to conduct a comprehensive evaluation of an entire enterprise architecture (i.e., an IT infrastructure of a large organization). The traditional models also do not enumerate and consider all of actual threat actions that have been observed in the wild. As no organization has unlimited resources to deploy every available protection, they must resort to a risk-based prioritization approach and deploy protections where they are needed the most or those with the biggest impact. To date, no commonly accepted methodology exists to allow organizations to look at actual threats in the wild and determine what kind of protection their cybersecurity architectures provide and where gaps exist [2].

Consequently, the decision makers, such as chief information security officers (CISOs), end up making their decisions on investments in cyber protections based on vendor recommendations and marketing pressure rather than using well-established risk management practices [2]. The threat-based approach to analysis of cybersecurity architectural capabilities allows organizations to consider the threat element in their risk management processes and make cybersecurity investment decisions informed by actual threats they are facing. Hence, there is an urgent need for a paradigm shift where organizations can look at their cybersecurity protections from the standpoint of an adversary to make threat informed risk decisions.

B. Problem Statement

As the concept of risk is a function of a security event or a scenario (i.e., threat exploiting a vulnerability), the probability of the event taking place, and the consequence of the event taking place [3], to fully exercise risk management practices, organizations need to factor in all these fundamental risk factors. No well documented and generally accepted methodology previously existed to allow for proper consideration of the threat factor in making risk-based decisions on investments in cybersecurity protections. This led to inadequate protections applied to organizational infrastructure, protection "blind spots", and wasted limited resources on protections that do not cover the actual threats organizations are facing or multiple protections covering the same limited threats.

The threat-based approach to evaluation of cybersecurity protections allows us to determine the best protection coverage against the actual cybersecurity threats organizations are facing. This is achieved by determining coverage of existing protections, identifying gaps (where threats without adequate protections exist), and overlaps (areas where multiple protections protect against the same types of threats thus unnecessarily multiplying costs) [2].

C. Organization of the Paper

There are five sections in this paper. Section I introduces the problem decision makers face when selecting cybersecurity capabilities for best protection against threats and provides an overview of the paper. Section II discusses the traditional threat modeling approaches and their deficiencies. Section III describes the threat-based approach to evaluation of cybersecurity architectures and associated protections as a novel approach to threat modeling. Section IV provides recommendations for future research, and Section V discusses expected findings and contributions.

D. Scope and Limitations

The cybersecurity protection (capability) categories are not standardized and differ widely from one vendor to another. Organizations are faced with thousands of cybersecurity products to choose from. In 2018, there were more than 1,200 cybersecurity vendors with approximately 6,000 products and more than 20,000 features [4]. While ideally, one would prefer to analyze all available cybersecurity products, including different models of the same product (e.g., Cisco ASA 5520 vs Cisco ASA 5550), such task would require tremendous resources and effort. To make the research manageable, this research will focus only on the major cybersecurity technology categories defined in Gartner's Magic Quadrant and Critical Capabilities [5].

II. Literature Review

A. Traditional Threat Modeling Approaches

Early attempts to formalize the threat modeling process for information systems were made by the Department of Defense in the late 1970s and early 1980s. One of the earliest dynamic threat analysis models was developed by ATT&T for the Strategic Defense Initiative as Security Vulnerability Analysis (SVA) for System Security Engineering (SSE) process. It was designed for structured enumeration of system security requirements through a ten-step process [6] and is known for the use of "threat logic trees" for threat decomposition.

A significant contribution to the development of threat modeling methodology was made by research sponsored by the National Security Agency and a group of researchers led by Bruce Schneier [7]. This model uses attack trees to visually represent possible threat actions and weigh them based on the risk, access, and cost to the adversary.

To date, the STRIDE methodology, a part of Microsoft's security development lifecycle (SDL) [8], is the most mature and widely used threat modeling methodology. It was developed by Microsoft Corporation in 1999 and named after major categories of threats occurring in the wild (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege) [9]. The STRIDE methodology enumerates possible threats and vulnerabilities in an information system, groups known threats into six categories, then describes various products and services each category applies to [9]. The system under consideration is deconstructed into components, and each component is analyzed for susceptibility to threats in each category which leads to discovery of associated vulnerabilities and assists with developing appropriate threat mitigation measures [10].

A similar threat modeling methodology based on STRIDE was developed by Gunnar Peterson [1]. The methodology is called DESIST, which stands for Dispute, Elevation of Privileges, Spoofing, Information Disclosure, Service Denial, and Tempering.

One of the more recent approaches to threat modeling is the Process for Attack Simulation and Threat Analysis or PASTA. What differentiates PASTA from other methodologies is the focus on business objectives as drivers for both information system requirements and associated security responses. The argument for the business objectives focus is based on the expectation that organizations in different industries face different types of threats and therefore only those impacting the organization should be mitigated [11].

B. Deficiencies of Traditional Models

What all traditional threat modeling methodologies discussed above have in common is that they do not consider the full spectrum of actual threat actions that have been observed in the wild. They enumerate a small subset of actual threat actions, and in some cases include theoretical or hypothetical threats. This approach works well on a small scale such as threat modeling for a specific data field (e.g., social security number records in a database), software application (e.g., during early stages of the development lifecycle), or a system component (e.g., cryptographic module of an authentication system).

The traditional modeling methodologies do not allow for a comprehensive, end-to-end evaluation of an entire enterprise and its cybersecurity architecture to determine what kind of protections the existing capabilities provide and where the gaps in need of decision makers' attention are. [2] This leaves the decision makers without a 'tool' to evaluate coverage (the level of protection) of individual cybersecurity capabilities (tools) and complex systems (architectures) they constitute.

C. Threat Modeling for Enterprise Cybersecurity Architecture

In 2015, the Department of Defense (DOD) introduced a new threat modeling methodology called NIPRNet SIPRNet Cyber Security Architecture Review (NSCSAR) (later renamed to DoDCAR) that allowed them to consider threats as a factor in the risk-based decision making process and, for the first time, look at the protection coverage of cybersecurity architectural capabilities from the standpoint of an adversary. This methodology continues to be widely used by DOD to identify gaps where protections do not exist and to inform the future investments into new protections. It also helps to identify protection overlaps (e.g., to inform decisions to retire redundant protections – use cases where two or more different products serve the same purpose and protect against the same type of threat). The Department of Homeland Security adopted this approach in 2018, and further improved it under the name .govCAR for the use by the federal agencies, other levels of government, and the public sector [12].

III. Cybersecurity Architecture Review

The main argument for the new approach to cybersecurity architecture review is in the need to enumerate and consider all stages and objectives of an attack, and associated threat

actions that have been observed being executed by adversaries in the wild.

A. Cyber Threat Framework

The comprehensive enumeration of threat actions is achieved using a cyber threat framework (CTF) which allows cybersecurity engineers to consider all previously used threat actions and create a common language to describe adversarial activities.

The approach is agnostic to a particular CTF, although it is most commonly used with National Security Agency's Technical Cyber Threat Framework (NSA/CSS NTCTF) v2.0 [14] and MITRE ATT&CK [13]. Both frameworks identify all different threat actions (sometimes called tactics, techniques, and protocols or TTPs) carried out by the adversaries in known cyber-attacks observed in the wild and group them by categories. For example, NTCTF 2.0, groups threat actions by breaking them down into six phases: Administration, Preparation, Engagement, Presence, Effect, and Ongoing process. Each phase then breaks down into two to five objectives – which generates a total of 21 objectives. Each objective can contain between two and twenty-one threat actions resulting in 186 individual threat actions. The inventory of threat actions is visually represented in a matrix.

C. Cybersecurity Capabilities, Flows, and Topologies

In the next step we identify the building blocks of target architectures: cybersecurity capabilities, their topologies (e.g., positions on the network), and network flows that are routed through those capabilities. The capabilities (also referred to as protections), defined as "combination of mutually reinforcing controls implemented by technical means, physical means, and procedural means […] typically selected to achieve a common information security or privacy purpose" [15] are vendor-agnostic representations of cybersecurity tools at an architectural level. Most frequently, capabilities represent technologies such as firewall or antivirus software (the methodology uses generic capabilities such as firewall instead of a particular vendor/model), but they can also represent non-materiel capabilities such as cybersecurity policies or NIST 800-53 controls [16].

D. Coverage Scoring and Analysis

Once selected, the architectural capabilities are arranged into a scoring matrix, with threat actions listed at the top as column headers and architectural capabilities on the left as row titles. Fig. 1 shows an excerpt from a sample scoring matrix with scores for one capability (e.g., firewall) and its three features against three threat actions (Inject database command, Leverage device swapping, and Send malicious email) in Delivery objective of the Engagement stage. If desired, capabilities can be further broken down into features (e.g., in the example above, the firewall capability was broken down into GeoIP blocking, Application filtering, and Protocol port enforcement [12].

Stage	Engagement								
Objective	Delivery								
Threat action	Inject database command			Leverage device swapping			Send malicious email		
Function	Protect	Detect	Respond	Protect	Detect	Respond	Protect	Detect	Respond
Capabilities									
Features									
Firewall	M	M	M	L	None	L	L	L	L
GeoIP Blocking	L	None	L	L	None	L	L	L	L
Application Filtering	M	M	M	L	None	L	L	L	L
Protocol Port Enforcement	L	L	M	L	None	L	L	L	L

Figure 1 - An excerpt from a sample scoring matrix

Each capability is scored for its ability to protect, detect, and respond to each threat action. This is achieved by answering the following questions at intersections of threat actions and corresponding capabilities (or features): a) can the capability (or feature) detect this threat action?; b) can the capability protect against this threat action?; and c) can the capability assist in recovery against this threat action? The answers are ranked on a scale from none, to some, moderate, or significant coverage. Detect, protect, and respond are three of five functions (Identify, Protect, Detect, Respond, and Recover) of NIST cybersecurity framework developed to provide a common language for describing cybersecurity risk among stakeholders [17].

DHS only uses the three functions in .govCAR analysis and further tailors their definitions to avoid ambiguity [12] as follows: The Protect function represents active measures with or without detection abilities that support the ability to limit or contain the impact of a threat action in cyber relevant time. The Detect function enables discovery of threat actions in cyber relevant time and require at least one sensor and an analytic function that operates on that sensor produced data. The Respond function provides data that support activities that occur after the threat actions have executed, including mitigation of the threat action or triggering further sensor data collection and analysis.

The functions are not mutually exclusive —a capability may be able to protect against a particular threat action but may not be able to detect or respond to the same threat action. For example, a firewall that is configured to drop all incoming traffic on port TCP/UDP:53 (a port typically reserved for standard DNS protocol) will protect against a threat action on this port but will not be able to detect nor log (respond to) activity associated with that threat action due to the traffic being dropped before such action can occur.

The answers to scoring questions form a capability coverage map – a visual representation of capability coverage, gaps, and overlaps against the threats. Coverage maps for multiple capabilities can then be overlayed on top of each other to evaluate the coverage of the entire organizational defense in depth architecture.

Figures 2-7 illustrate the threat modeling results for six generic cybersecurity capabilities and Fig. 8 illustrates their combined effects. The capability coverage maps are color coded as defined in Fig 9. The capability in Fig. 2 has limited coverage against 20 threat actions and no significant of moderate coverage. The capability in Fig. 3 has limited coverage against 28 and significant coverage against 6 threat actions. The capability in Fig. 4 shows limited coverage against 76 and moderate against 1 threat action. The capability in Fig. 5 has limited coverage against 155, moderate coverage against 61,

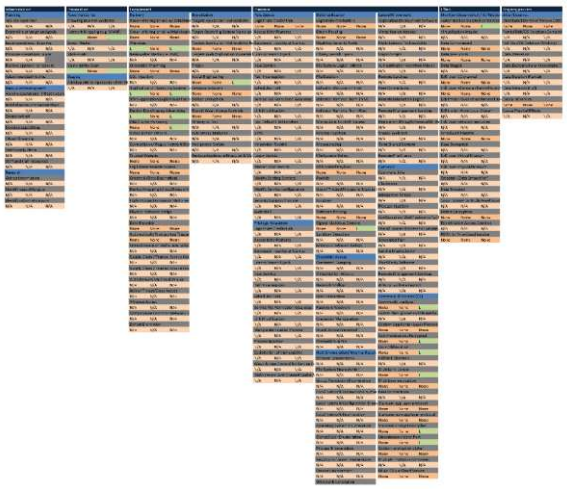

Figure 2 - Capability A Coverage Map

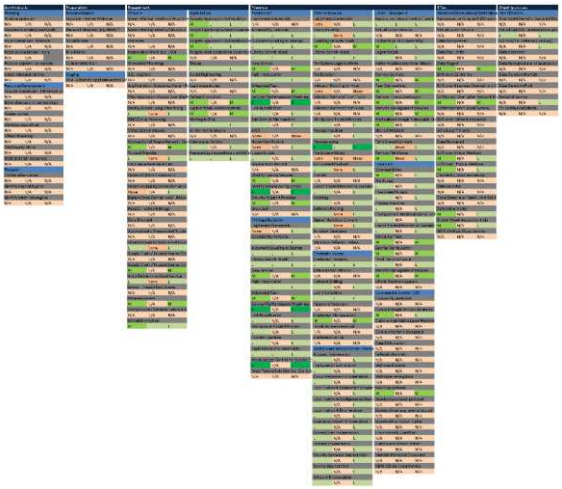

Figure 3 - Capability C Coverage Map

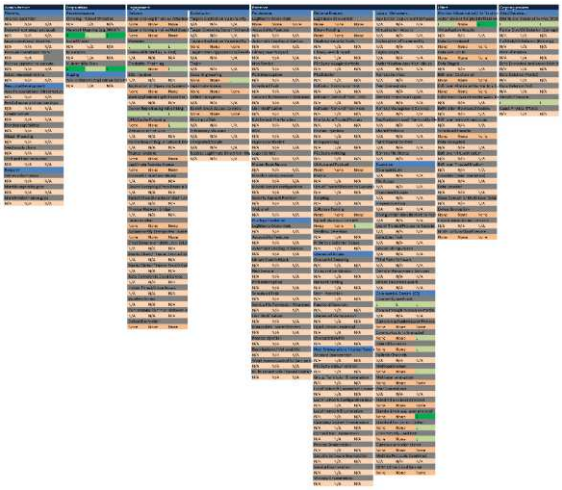

Figure 4 - Capability B Coverage Map

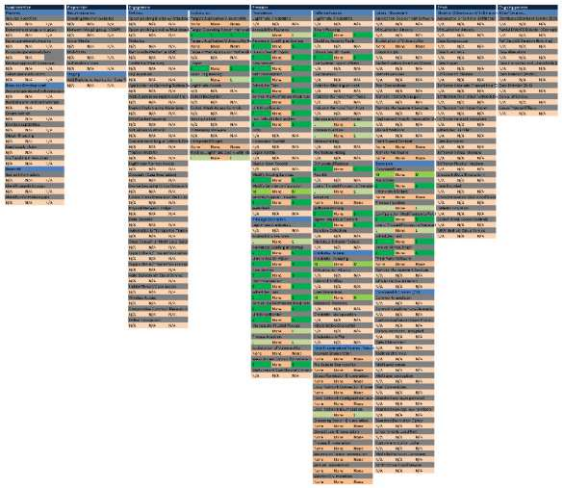

Figure 5 - Capability D Coverage Map

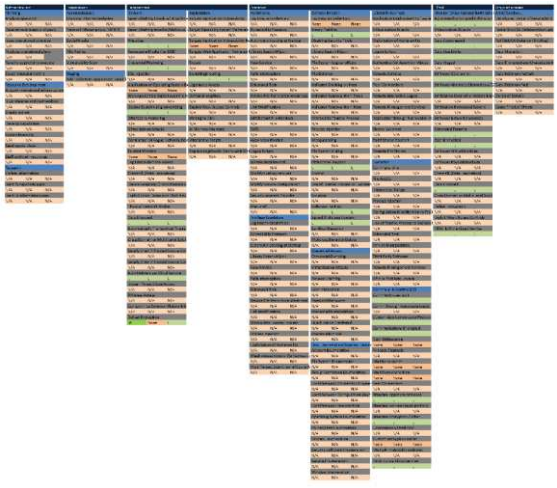

Figure 6 - Capability C Coverage Map

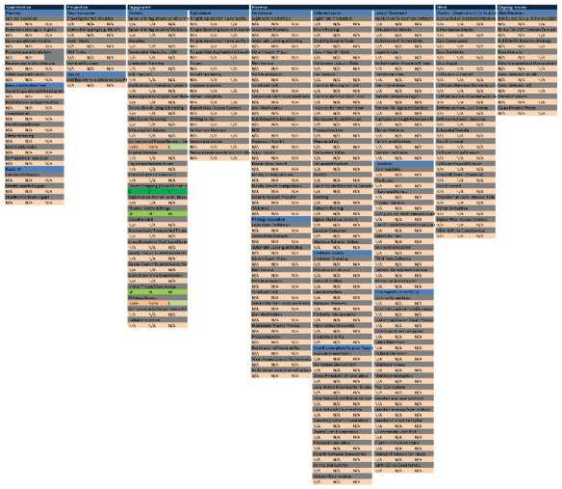

Figure 7 - Capability E Coverage Map

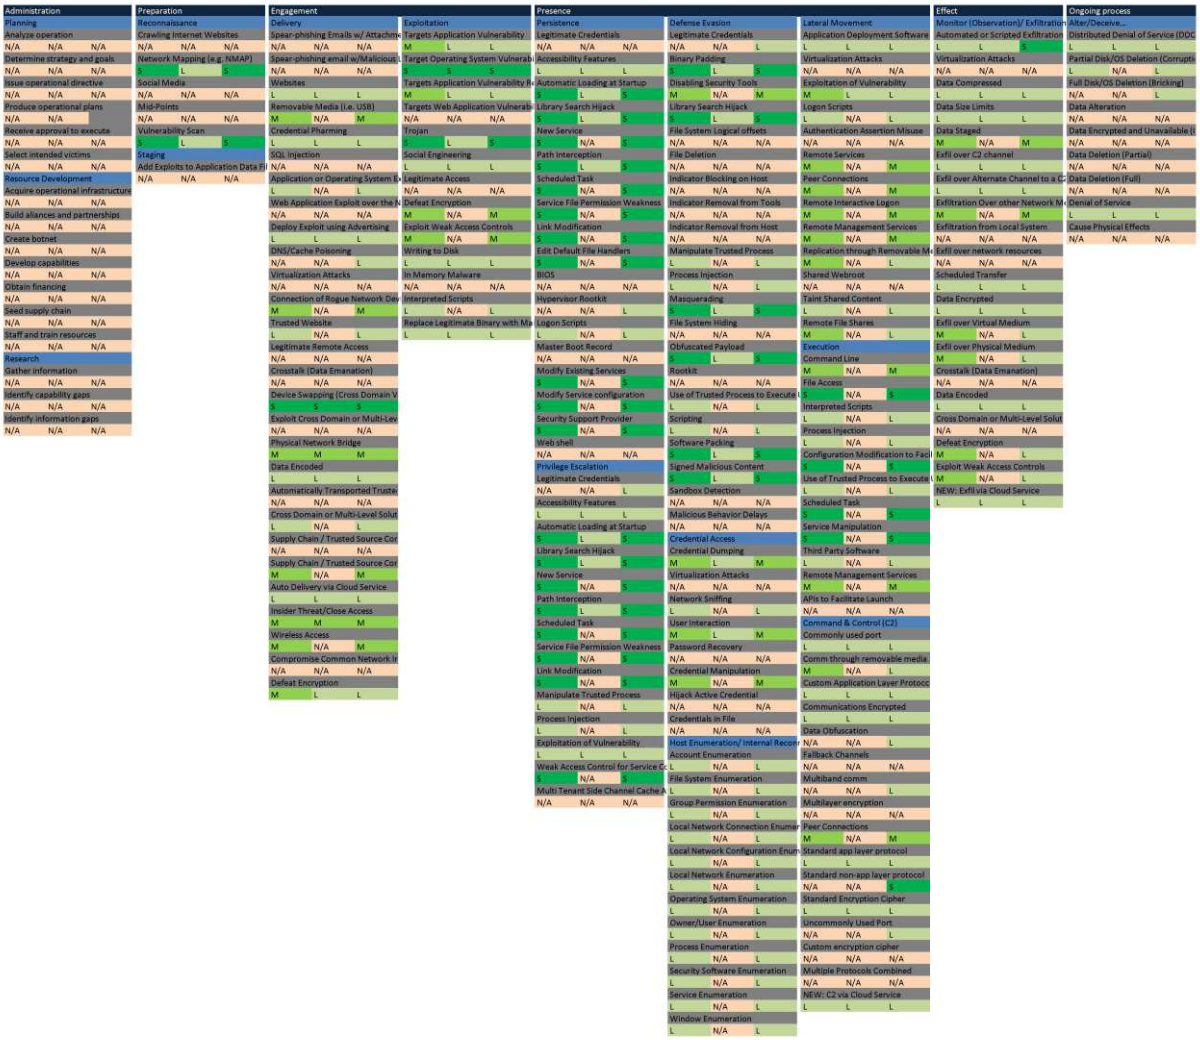

Figure 8 - Coverage Map Overlay for Six Capabilities (A-F)

and significant coverage against 10 threat actions. The capability in Fig. 6 has limited coverage against 24, moderate coverage against 8, and significant coverage against 59 threat actions. Finally, the capability in Fig. 7 has limited coverage against 2, moderate coverage against 6, and significant coverage against 3 threat actions. In Fig. 8, we see the results of combining (overlaying) the coverage of six individual capabilities in an enterprise architecture. The overall increase in coverage is evident from the change of matrix colors from predominantly coral pink to green. The combined capabilities have limited coverage against 305, moderate coverage against 76, and significant coverage against 78 threat actions.

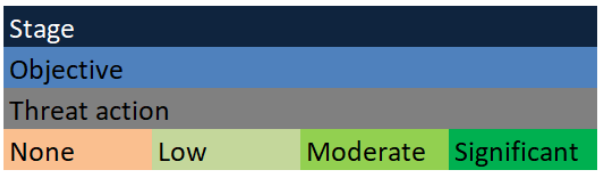

Figure 9 - Coverage Map Color Codes

The threat actions can also be evaluated based on their prevalence (frequency of occurrence in the wild) and maneuverability (the number of different threat actions that can be used to achieve the same objective) with results visually represented on a threat heat map. The heat map can be overlaid on top of any coverage map to better understand and prioritize future protections focus.

IV. Recommendations for Future Research

To fully reach the potential of the cybersecurity architecture review threat modeling methodology we propose a development of a decision-making framework for enhancing cybersecurity capability portfolios to maximize protect, detect, and respond coverage against cyber threat actions.

In order to achieve the research goal, we have identified five research questions along with input data (and its sources), research methods, and output data. Fig. 10 outlines the research questions and provides a high-level overview of the relationships between each research question, data inputs,

Research Question	Input Data	Method	Output Data	How Does it Answer the Goal?
RQ1: To what extent are organizations protected against known cybersecurity threats?	CTF Threat heat maps Common cybersecurity technologies	Directed literature review white papers. Multicriteria decision analysis of capabilities against CTF.	Coverage maps	Identifies how common capabilities protect actual threats Identifies gaps and overlaps
RQ2: Do organizations make the most efficient investment decisions that maximize capability protect/detect/respond coverage against the actual cybersecurity threats?	Coverage maps. Common cybersecurity technologies	Survey of a representative sample of organizations Interviews with select professional groups	Capabilities in use by surveyed organizations Coverage maps for surveyed organizations	Identifies how capabilities used by surveyed organizations protect threats Identifies common gaps in capabilities at surveyed organizations Contrasts actual vs. ideal portfolio coverage maps
RQ3: Is the selection of capabilities by organizations influenced by vendors or the actual threat landscape?	Capabilities in use by surveyed organizations Coverage maps for surveyed organizations Vendor market share data	Comparative analysis of coverage maps Statistical analysis to compare coverage vs market share	Distribution of capabilities between actual threat actions and major vendors	Determines the major drives behind the capability types and recommendations for optimization.
RQ4: Do different demographics (CISO vs network administrator, female vs male) or industries (government vs energy vs education) perceive the capability coverage differently and how?	Coverage maps for surveyed organizations Demographics and industry data	Interviews with select professional groups to determine their perception of capability coverage.	Perceived coverage maps for select demographics.	Determines other factors that may influence selection of capabilities
RQ5: To what extent do the major cybersecurity capabilities deployed at surveyed organizations overlap?	Coverage maps. Capabilities in use by surveyed organizations Coverage maps for surveyed organizations	Qualitative and quantitative analysis of collected data	Coverage overlap maps	Determines overlaps – areas where multiple capabilities protect against the same threat actions.

Figure 10 - Proposed research questions, data inputs, methods, and data outputs

methods to be used in analysis, data outputs (resulting data), and the relationships of each research question to the research goal.

V. Expected Findings and Contribution

A. Expected Findings

The recommended future research is expected to confirm the hypothesis that cybersecurity investment decisions are not threat driven but rather based on the market and vendor pressure. We also expect to demonstrate that leads to organizations having significant gaps in protections against known threat actions.

B. Expected Contribution

We seek to demonstrate that the current approach to selection of cybersecurity architectural capabilities is inadequate and expect to provide cybersecurity practitioners with a better decision-making framework for enhancing cybersecurity capability portfolios to maximize protect, detect, and respond coverage against the current cybersecurity threat actions.

References

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

[2] Bokan, B., & Santos, J. (2021). Managing Cybersecurity Risk Using Threat Based Methodology for Evaluation of Cybersecurity Architectures. Systems and Information Engineering Design Symposium (SIEDS) (pp. 1-6). IEEE.

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

[4] N. Miller, "With More Than 1,200 Cybersecurity Vendors in the Industry, How Do You Stand Out?," 8 May 2018. [Online]. Available: https://www.mcafee.com/blogs/enterprise/with-more-than-1200-cybersecurity-vendors-in-the-industry-how-do-you-stand-out/.

[5] Gartner, "Gartner Magic Quadrant & Critical Capabilities," 13 May 2020. [Online]. Available: https://www.gartner.com/en/research/magic-quadrant.

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

[7] 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.

[8] 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.

[9] 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/.

[10] 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.

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

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

[13] The MITRE Corporation, "Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK)," 17 February 2019. [Online]. Available: https://attack.mitre.org/.

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

[15] 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.

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

[17] National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity Version 1.1," Gaithersburg, 2018.