Standardizing Zero Trust Implementation Using NIST SP 800-53 and NIST SP 800-207

A. Gupta[1], C. Pasciuto[2], K. Lai[3], L. Sayre[4], Y. Xu[5]
Heinz College of Information Systems and Public Policy
Carnegie Mellon University
Pittsburgh, USA
[1]aarushg2@andrew.cmu.edu, [2]cpasciut@andrew.cmu.edu, [3]klai2@andrew.cmu.edu,
[4]lsayre@andrew.cmu.edu, [5]yunhanx@andrew.cmu.edu

***Abstract*—Through conducting research on existing security standards and Zero Trust frameworks, the Carnegie Mellon University Heinz College capstone team, in conjunction with the Software Engineering Institute, developed a mapping between security controls and Zero Trust maturity rankings to provide preliminary guidance on implementing Zero Trust within private and public enterprises. The team's mapping work was supplemented by a sensitivity analysis to determine term frequency in mapping outcomes, as well as qualitative feedback in the form of survey responses from security professionals and interviews with Zero Trust leaders.**

***Index Terms*—Zero Trust, cybersecurity, maturity, NIST SP 800-53, NIST SP 800-207, CISA Zero Trust Maturity Model Draft, sensitivity analysis**

Executive Summary

The variability in Zero Trust (ZT) definitions across organizations and industries makes achieving ZT a challenge, as professionals seek to ensure secure technology environments in a world with ever-increasing internal and external threats. Developing a method to implement and measure the effectiveness of an enterprise's Zero Trust Architecture as it pertains to security controls and foundational ZT concepts is the first step in moving ZT from a buzzword to a standard practice. To establish a preliminary framework through which organizations can begin to implement ZT, this team focused on identifying the extent to which the National Institute of Standard and Technology (NIST) Special Publication (SP) 800-53 controls and enhancements promote ZT, in alignment with NIST SP 800-207 and the Cybersecurity & Infrastructure Security Agency (CISA) Zero Trust Maturity Model Draft. Understanding how varying levels of maturity can be achieved through the implementation of specific controls will encourage organizations to migrate from traditional information security practices to more rigid and effective ZT procedures, ensuring business continuity and customer/employee confidentiality.

I. Background

A. Introduction to Zero Trust

NIST SP 800-207 defines Zero Trust as "the term for an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources," with a Zero Trust Architecture (ZTA) defined as a security design that is reliant on ZT principles "to plan industrial and enterprise infrastructure and workflows". While traditional cybersecurity principles focus on the use of perimeter-based defenses, ZT assumes that attackers and threats already exist within an enterprise environment; because of this, "there is no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the internet) or based on asset ownership (enterprise or personally owned)," and risks to assets and business functions must be continuously assessed and analyzed [1]. By assuming that "there is no traditional network edge," ZT stipulates that all users both internal and external to an organization's network must be continuously authenticated and validated before being granted access to applications, data, or any other organization-level resource [2]. Through this enforcement of "accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised," ZT provides a security model that is adept at addressing the risks associated with remote workforces, cloud infrastructure, and other potential threats that exist across the modern technological enterprise [1].

NIST SP 800-207 prescribes seven core tenets to which ZTA must adhere:

1) All data sources and computing services are considered resources.
2) All communication is secured regardless of network location.

3) Access to individual enterprise resources is granted on a per-session basis.
4) Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes.
5) The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6) All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
7) The enterprise collects as much information as possible about the current state of assets, network infrastructure and communications and uses it to improve its security posture.

II. Objective and Scope

A. Project Justification

While the concept of ZT has been around for several years, recent high-profile cybersecurity incidents, such as the 2020 SolarWinds hack and the 2021 Colonial Pipeline attack, have brought the topic of ZT back into the spotlight as a more secure solution for combating an evolving and increasingly complex threat landscape. The May 2021 Cybersecurity Executive Order, issued in response to these incidents, mandated that all Federal Civilian Executive Board (FCEB) agencies develop and submit formal plans to implement ZTA across their infrastructure. As ZT adoption accelerates in the coming years, it is likely that both public and private organizations will seek guidance and recommendations on how to best implement ZTA within their systems.

As one of the nation's premier federally funded research and development centers, the Carnegie Mellon Software Engineering Institute (SEI) has historically spearheaded the generation of new cybersecurity technologies and guidance for both government and private sector organizations to adopt and reference. Given the rise in prominence of ZT, the SEI has the opportunity to develop a key reference for government and private enterprises to consider as they migrate to ZTA. Such a product can facilitate and streamline the successful implementation of ZT security models across organizations in the United States, thus minimizing the time and resources required to implement ZT controls, the occurrence of future cybersecurity incidents, and ultimately, the negative economic and reputational impacts suffered by companies as a result of an incident.

B. Objective

This project aims to create the foundations for a formalized ZT framework that can aid organizations in the selection and implementation of NIST SP 800-53 controls, or similar controls from a related framework. The NIST SP 800-53 mapping, the cornerstone of this capstone project, intends to link these controls with the CISA ZT Maturity Model and the guidance offered by NIST SP 800-207, providing organizations with the ability to identify controls and control families where they are deficient, in regards to ZTA implementation. Through this identification process, organizations should be able to self-assess their own level of maturity with regards to migrating to ZTA, and determine which controls and control families should be prioritized by their security organization. The Sensitivity Analysis and the Web Application are meant to enhance the NIST SP 800-53 mapping by providing organizations with additional tools and resources to better understand the content described in each Maturity Model pillar, and the ability to quickly procure and customize a list of NIST SP 800-53 controls that is relevant to their needs and interests as they relate to ZTA implementation. At a high level, summary statistics from the Control Decision Survey can further inform organizations on how different companies and verticals are considering ZT strategies and relevant controls.

C. Project Description

The project team created four distinct products that organizations can reference to assist in their migration to ZT:

1) *Control Decision Survey & Interviews* - The project team conducted a survey of security managers across different companies and industries on their respective selection and implementation of NIST SP 800-53 controls as they relate to ZT. The purpose of the survey is to understand how different organizations prioritize control selection as they implement ZTA, as well as develop an overview of how widely adopted ZT principles currently are across separate verticals.
2) *NIST SP 800-53 Control Mapping* - The project team mapped over 1000 NIST 800-53 controls to each of the five pillars identified in the CISA Zero Trust Maturity Model Draft; each control was categorized into the pillar(s) the control related to, designated a primary and secondary function (where necessary), and assigned a ZT maturity rating in accordance with the CISA document. The purpose of this mapping is to synthesize NIST

SP 800-53 controls with the ZT maturity ratings defined by CISA, enabling organizations the ability to self-assess the sophistication of the controls and control families inherent in their system to determine future areas of focus for ZTA implementation.

3) *Sensitivity Analysis* - The project team implemented a complete sensitivity analysis of the CISA Zero Trust Maturity Model text; each of the five pillars and its corresponding functions were reviewed to determine the frequency of critical terms that appear under each. The purpose of this analysis is to identify the key content of each pillar and its functions, enabling organizations to quickly understand they relate to the three maturity ratings introduced by the model.
4) *Web Application* - The project team created a web application to accompany the results of the NIST SP 800-53 control mapping. The purpose of this web application is to provide organizations with the ability to quickly procure a list of 800-53 controls and control families that is customized to their organizational ZT interests and needs.

The "Deliverables" section of this report discusses the methodologies utilized to complete each deliverable in further detail, while the conclusion outlines next steps and future work that can be conducted on each product outlined above.

III. Business Impact

While some organizations may have already considered implementing aspects of ZTA into their current information security systems, the majority of individuals and entities across a variety of industries interviewed for this capstone project indicated that serious considerations for transitioning to ZTA remain in the preliminary stages. The research executed by this capstone project team is intended to assist companies intent on transitioning to a ZT model in the selection and implementation of relevant NIST SP 800-53 controls. By facilitating this selection of controls, this research can help streamline the successful implementation of ZT security models across organizations in the United States. Below are three of the positive impacts this project team aims to achieve:

- *Minimization of Time & Resources* - Organizations can utilize the framework provided by this capstone project to efficiently select NIST SP 800-53 controls for implementation; using the maturity ratings and primary functions assigned to each control throughout the mapping process, organizations can identify which CISA ZT Maturity Model pillars they are deficient in without having to devote additional time and resources into investigating. For these pillars, organizations have the ability to view, select, and prioritize relevant control enhancements for implementation to elevate their ZT maturity.
- *Minimization of Future Cybersecurity Incidents* - By utilizing this framework to assist in transitioning to ZTA, organizations can improve their overall cybersecurity postures to better address the current threat landscape. According to IBM's "Cost of a Data Breach" Report, average data breach costs totaled $4.24 million in 2021, with compromised credentials identified as the most common attack vector for data breaches. The report finds that the "average cost of a breach was USD 1.76 million less at organizations with a mature zero trust approach, compared to organizations without zero trust" [3].
- *Minimization of Residual Economic Reputational Impacts Caused by Incidents* - By improving their threat security posture through the implementation of ZT principles, organizations can better protect themselves from the economic and reputational harm associated with high profile cybersecurity incidents. Residual costs can manifest itself in many ways for victims of data breaches; legal fees from lawsuits, operational downtime, regulatory fines, and revenue loss from reputational harm can financially cripple an organization. In the infamous case of the 2017 Equifax breach, Equifax lost an estimated $4 billion in stock market value within a week of the incident's announcement, an additional $439 million in costs directly associated with the breach by the end of 2017, and was ordered by US courts to spend $1 billion in cybersecurity enhancements under court supervision [4].

IV. Deliverables

A. Control Decision Survey and Interviews

To understand the current landscape of ZT, the team generated a survey and conducted various interviews. This allowed the team to understand how security professionals are currently implementing ZT across industries. The purpose of the surveys was to develop perspective and compile quantitative data concerning security professionals' view of ZT progression among a diverse range of organizations and verticals. In addition, the interviews provided the team with more detailed, qualitative information that fostered a more comprehensive understanding of security environments across organizations and of ZT leaders' efforts in implementing ZT.

1) Methodology: The team initially drafted a working set of a dozen or so questions that focused on how organizations are currently approaching their transitions to ZTA, with specific consideration given to how companies assess their own strengths and weaknesses across the NIST SP 800-53 control families that are critical to the pillars and functions identified in the CISA Zero Trust Maturity Model. The draft was then submitted to both the project team advisor and the client for feedback, and revisions were made to several of the questions' phrasing and the rating mechanisms. Once the list of questions were finalized and approved, a survey was created using the Qualtrics software platform (the complete list of survey questions featured in the Qualtrics survey is included below). Links to the survey were then sent out by individual team members to security professionals across their networks; additionally, the team coordinated with Heinz College Career Services to distribute the survey to alumni from the Information Security Policy & Management program for their consideration as well. For the interview portion, the team constructed a list of various professionals to identify the perspectives of security leaders, security engineers, and Zero trust experts. To capture a wide variety of perspectives, the team interviewed several individuals from different organizations and industries, including:

- CISO, major financial institution
- CISO, artificial intelligence company
- Networking Lead, major financial institution
- Zero-Trust SME, insurance institution
- Zero-Trust Researcher, private and public sector
- Cloud Zero Trust Researcher, private sector and cloud securityboard
- Retired US General and US Cybersecurity Leader, FFRDC

(Note: The identities of selected interviewees have been masked to protect their privacy and organization confidentiality).

Survey Questions

- What is your organization's industry?
- Across your industry, how widely adopted are ZT principles?
- Please rank your consideration of the following aspects when implementing ZT (from 1-5, 1 being most important): Critical Assets, Business Goals/Objectives, Risk Assessment, ROI Analysis, Critical Assets, Capability Maturity
- Do you consider other aspects when implementing ZT?
- (If applicable) Of the NIST SP 800-53 control families, which control family have you found to be most critical and pertinent to your organization in terms of your ZTA?
- (Optional) Of the NIST SP 800-53 control families, in which control family do you believe your organization is most deficient currently, with regards to ZTA migration?
- Of the NIST SP 800-53 control families, which control family do you believe your organization excels at the most, with regards to ZTA migration?
- What leads you to believe your organization excels at your option of choice at the most?
- Which of the following ZT maturity pillars does your organization place the strongest emphasis on?
- Why does your organization place the strongest emphasis on it?
- On a scale of 1-5, how soon could your organization begin migrating to a ZTA?
- What tools would facilitate your decisions in choosing certain controls to support ZT migration for your organization?
- Which tools would facilitate the measurement of effectiveness of controls to support ZT at your organization?

2) Findings and Analysis: The 35 survey respondents come from 12 different industries: National Security, Aerospace/Defense, Government, Software as a Service, Technology, Telecommunications, Manufacturing, Pharmacy, Insurance, Legal, Financial, and Entertainment. To summarize the findings, the team grouped the survey questions into five buckets:

1. Across your industry, how widely adopted are Zero Trust principles?

On a scale of 1 (not adopted or considered) to 4 (widely adopted), the average response was 2.5, between minimally adopted and relatively adopted. However, as the respondents represented are more likely to adopt high-level security measures since they are security professionals, it can be surmised that the population average is lower than 2.5.

2. What other things do you consider when implementing ZT?

The top five considerations from high to low are critical assets, ROI analysis, business goals/objectives, risk assessment, and capability maturity. Respondents also mentioned other considerations: timeline, scalability, service impact, employee experience, user interruption/usability, implementation cost, and competitive advantage.

3. Which control families do you consider to be the most important, with regards to ZTA migration?

The three most critical and pertinent control families from respondents are IA (Identification and Authenti-

cation), AC (Access Control), and CM (Configuration Management). The most mentioned control family that respondent organizations excel at the most is AC (Access Control).

4. On which control pillars does your organization place the strongest emphasis?

The Identity control pillar received the most votes, as respondents frequently agreed that "identity is the new perimeter" and "access to resources highly depends on the user and device accessing." Network/Environment is considered "more manageable" than Device and Data. Some respondents also expressed that "organizations are focused on securing their data to achieve higher confidence in their services from consumers."

5. How soon could your organization begin migrating to a ZTA?

On a scale of 1 (high maturity already implemented) to 5 (require 5+ years), the average is 2.8, between require ¡1 year and require 1-2 years. The average is lower than expected, likely due to two reasons: First, Zero Trust remains a popular "buzzword" and relatively novel concept, which leaves the term subject to a variety of different interpretations; this could result in respondents misunderstanding what a complete ZTA entails. Second, someone who has a background in ZT could be more likely to fill out the survey and is, therefore, more confident in assessing their organization's ability to implement ZTA.

The interviews with security experts also yielded useful insight regarding the current landscape of the security industry, as well as the maturity of ZT implementation. There were four common themes that stood out to the team during these discussions:

1. Managing Identity is one of the most significant concerns

During the interviews, the team consistently heard that identity was the pillar of most concern. Security professionals stated that this was a significant concern due to the complex environments they are managing. Professionals are worried about how they will be able to achieve an optimal maturity rating in this pillar in order to fully leverage a ZTA.

2. Technological debt hinders ZT adoption

Many professionals and organizations stated that their technological debt was the main hindrance to their adoption of ZTA. The interviews uncovered that many organizations do not have the budget, staffing, or resources to rewrite or re-architect all their legacy systems. Instead, companies have focused on implementing ZT where they can, but organizations do not have fully developed plans to bring their legacy systems into a ZTA. In these interviews, it was often stated that if an organization has less technical debt, it would be able to accelerate its ZT transformation.

3. The industry needs a standardized way to measure control effectiveness

From the interviews, the team discovered that many organizations struggle with a formal way to measure the effectiveness of their controls. Companies adopt a wide variety of methods to measure the effectiveness of their implemented controls. However, in order to effectively measure the maturity of their ZT transformation, companies need a standardized approach to measuring the effectiveness of their controls.

4. The industry lacks prescriptive guidance

The most resounding feedback heard during the interviews was that the industry needs definitive, prescriptive guidance for ZT implementation. Our interviewees often mentioned that the ZT guidance is too high level. Industry leaders and technical experts need guidance on where to start implementing ZT and a framework that helps them to prioritize and implement controls to accelerate their ZT transformation.

B. NIST SP 800-53 Mapping

The purpose of mapping NIST SP 800-53 controls to the CISA Zero Trust Maturity Model is to provide organizations that utilize the NIST Risk Management Framework (or a similar framework) with transparency on how these controls relate to significant ZT functions and capabilities. In an ideal scenario, organizations can reference this mapping to identity areas in which they are most deficient in, with regards to transitioning towards ZTA, and utilize this knowledge towards the prioritization and implementation of controls and control families most critical to their needs.

1) Methodology: In Fall 2021, CISA released the ZT Maturity Model in draft format. The model is intended to aid FCEB agencies in their transitions to ZTA, as mandated by Executive Order 14028, by providing organizations with the ability to self-assess their current information systems and security policies as they relate to ZT principles reflected in NIST SP 800-207. The model organizes over 30+ functions under five primary pillars that form the foundation of ZTA:

1. Identity
2. Device
3. Network/Environment
4. Application Workload
5. Data

Examples of functions discussed in the ZT Maturity Model include: Authentication, Risk Assessment, Threat Protection, Access Authorization, Compliance Monitoring, and Governance Capability (please reference the

CISA Zero Trust Maturity Model for a comprehensive list of all functions and their descriptions). The model introduces a three distinct maturity levels to categorize each function; a summary description of each rating is outlined below:

Traditional: Functions that feature manual configurations and assignment of attributes, or static security policies; proprietary and inflexible pillars of policy enforcement, manual incident response and mitigation deployment.

Advanced: Functions that feature some cross-pillar coordination, centralized visibility identity control, policy enforcement based on cross-pillar inputs and outputs, some incident response to pre-defined mitigations, increased detail in dependencies with external systems, and/or some least privilege changes based on posture assessments. (Note: "Advanced" is the baseline rating needed to achieve a level of ZT compliance under the CISA Maturity Model).

Optimal: Functions that feature fully-automated assigning of attributes to assets and resources, and/or dynamic policies based on automated/observed triggers.

To start the mapping process, the team exported all 1300+ NIST SP 800-53 controls into a Google Spreadsheet. It then assigned 2-3 team members to each control family to begin categorizing and mapping the controls to the pillars, functions, and rating scheme outlined in the CISA Maturity Model. Team members followed these general steps in their approach to mapping each NIST SP 800-53 control:

1) Review of the Control Name, Control Text, and Discussion
2) Selection of one or more pillars to which the control relates
3) Assignment of a "Primary Function" to the control
4) Assignment of a "Secondary Function" to the control (where applicable)
5) Selection of one of three "ZT Control Maturity Ratings"

Controls that contained text or discussion that was either irrelevant or contradictory towards ZT principles were highlighted in yellow by team members for further review, with notes added to the Comment section about why the control was highlighted. Once a preliminary analysis of all NIST SP 800-53 control families was conducted, team members were then assigned to review a different set of control families than the ones they initially assessed. For this second phase of the mapping process, team members followed these general steps in their approach:

1) Controls were highlighted in green if the team member agreed with the mapping
2) Controls were highlighted in red if the team member disagreed with any portion of the control mapping; this could include improper assignment to a pillar, or incorrect selection of primary/secondary functions

   NOTE: Controls that were highlighted in yellow were left alone until phase three of the mapping process
3) Controls that were highlighted in red were brought forth during weekly team meetings for further discussion; for several weeks, the team regularly spent 20-30 minutes at the conclusion of each team meeting reviewing red controls and coming to an agreement on how to best map and categorize the control in question. By the end of this second phase, all controls were highlighted in either green or yellow, with no red controls remaining.

For the final phase of the mapping process, the team assigned members to different control families to review all controls that had previously been highlighted in yellow. Team members brought forth highlighted controls that they deemed relevant for further discussion during weekly team meetings, and the team collectively worked together to map these controls and highlight them in green once agreement had been established on their analysis. Controls that remained highlighted in yellow were then organized into a separate document for further research and analysis; in general, these highlighted controls fit several themes that contradicted tenets of ZT outlined in the CISA Maturity Model and NIST SP 800-207. The following section includes summary findings on the potential conflicts these controls have with ZTA implementation, and recommendations on how future versions of NIST SP 800-53 may mitigate these possible friction points.

2) Findings and Analysis: A number of mapping ambiguities arose as the team established alignment between NIST SP 800-53 controls and the chosen Zero Trust frameworks. Initially, it found that some controls and even entire control families posed little to no relevance to ZT; however, over the course of its research and mapping efforts, the team discovered several topics addressed in NIST SP 800-53 that seem to conflict with ZT fundamentals altogether. This realization prompted the team to perform an analysis on key areas within NIST SP 800-53 that warrant closer attention in order to better champion ZT.

(Note: The controls and enhancements provided as examples below are not comprehensive lists. Rather, they are provided as examples of overarching themes the team uncovered in its mapping efforts. Further commentary can be found within the team's mapping document).

Methodology: Gap Remediation

Remediation of mapping conflicts occurred simultaneously with the mapping process. As the team discovered controls that seemed to contradict or were inapplicable to ZT Migration, it highlighted the controls and documented areas of concern pertaining to each control and its ZT mapping. Afterwards, the team discussed the consistencies across the highlighted controls. Rather than analyzing controls and their enhancements on an individual level, the team grouped controls into various "gap" topics.

Results

The primary gap topics the team determined are Shared Systems and Account Environments, Privileged User Accounts and Access, Exceptions, and Boundaries. A final topic "Irrelevant Controls" is included to address the controls within NIST SP 800-53 that do not maintain relevance to ZT.

1. Shared System/Account Environments

To truly achieve an optimal ZT maturity ranking, organizations must retire shared accounts. Though efficient, shared systems and accounts introduce greater security risk to an organization. If a fully established ZTA requires dynamic authentication and authorization of users and devices (see NIST SP 800-207), shared systems and accounts are incapable of enforcing ZT. Shared accounts do not authenticate at the user or device level, instead granting users access so long as they possess the shared password.

For organizations to achieve an optimal ZT maturity level, they must transition from static to dynamic account access, which inherently excludes the use of shared accounts. A few controls within NIST SP 800-53 that the team found to contradict this include:

- CA 6(1) - Authorization — Joint Authorization – Intra-organization
- CA 6(2) - Authorization — Joint Authorization – Inter-organization
- IA 2(1) - Identification and Authentication (organizational Users) — Multi-factor Authentication to Privileged Accounts
- IA 2(5) - Identification and Authentication (organizational Users) — Individual Authentication with Group Authentication

Although these controls/enhancements may present measures through which security risks can be limited for shared accounts, proposing the use of shared accounts at all is contradictory to ZT.

2. Privileged User Accounts/Access

While ZT methodology stresses the importance of limiting privileges and permissions across organizations, it is infeasible to expect organizations to remove all administrative and privileged credentials from their systems; many of the controls highlighted in this section require privileged accounts for specific purposes, such as software installation and vulnerability scanning actions. Although the current version of NIST SP 800-53 does contain controls that reflect critical principles of effective Privileged Access Management, the current revision of the framework falls short of addressing the ZT topics of Just-in-Time and Just-Enough access. Just-in-Time and Just-Enough access replace the traditional concept of having standing administrative privileges by instead granting privileged access on a case-by-case basis. Requests must "mee[t] a set of criteria indicating it's a standard one," and the requester is "given access to only what they need to complete the job – for a specific length of time necessary to complete the task. When the user has finished the task, they lose their access privileges" [5]. As seen in incidents such as SolarWinds and the Colonial Pipeline, compromised credentials can enable threat actors to quickly gain unbounded access to an organization's entire infrastructure; as such, removal of access to standing privileged credentials can mitigate many of the threats associated with modern cybersecurity challenges. By including controls that specifically address these concerns, NIST SP 800-53 can better align with both today's evolving threat landscape and many tenets of ZT.

Several NIST SP 800-53 controls that fall into this category include:

- AC 2(7) - Account Management — Privileged User Accounts
- AC 3(7) - Access Enforcement — Role-based Access Control
- AC 6(5) - Least Privilege — Privileged Accounts
- AC 6(8) - Least Privilege — Privilege Levels for Code Execution
- RA 5(5) - Vulnerability Monitoring and Scanning — Privileged Access

3. Exceptions

Although it may not be feasible to deny all exception requests – as this, many times, limits an organization's ability to effectively carry out operations – the concept of exceptions does not support ZT, let alone basic information security. Controls that establish proper actions in the case of exceptions are useful, but to meet optimal expectations as far as ZT concerns itself demands the most rigid response from security organizations.

Controls and enhancements that the team found to encourage exceptions (along with proper standard security controls) include:

- AC 2(9) - Account Management — Restrictions on

Use of Shared and Group Accounts
- AC 3(10) - Access Enforcement — Audited Override of Access Control Mechanisms
- CM 7(4) - Least Functionality — Unauthorized Software – Deny-by-exception
- PT 6(1) - System of Records Notice — Routine Uses
- PT 6(2) - System of Records Notice — Exemption Rules

4. Boundaries

Boundary protection is still relevant in ZT; however, how boundaries are viewed in a ZTA differs significantly from a traditional security model. As the National Security Agency (NSA) states, "Zero Trust is a security model, a set of system design principles, and a coordinated cybersecurity and system management strategy based on an acknowledgment that threats exist both inside and outside traditional network boundaries" [6]. The mindset differs from a more traditional security model that focuses on the external boundaries of organizations. The controls above reflect on that traditional mindset that focuses on an outward-facing threat. ZT changes the concepts of boundaries by identifying that threats both exist within and outside of organizations. In this way, ZT shifts the boundaries down to a much more granular level, and some would even argue that individuals are the boundaries in a ZTA. In essence, the threat to an organization is omnipresent throughout the network. The controls above do help to support ZT and a secure architecture; however, they need to be expanded upon and redefined to account for a much larger definition of boundary protections that truly support ZT.

A majority of the gap controls related to boundary protection comprised the System and Communications Protection (SC) control family within NIST SP 800-53, e.g.,

- SC-7(3) - Boundary Protection — Access Points
- SC-7(4) - Boundary Protection — External Telecommunications Services
- SC-7(5) - Boundary Protection — Deny by Default – Allow by Exception
- SC-7(7) - Boundary Protection — Split Tunneling for Remote Devices
- SC-7(8) - Boundary Protection — Route Traffic to Authenticated Proxy Servers
- SC-7(9) - Boundary Protection — Restrict Threatening Outgoing Communications Traffic

5. Irrelevant Controls

Examples of control families and controls that the team deemed entirely irrelevant to the concept and implementation of ZT include:

- PM-19 - Privacy Program Leadership Role
- PM-24 - Data Integrity Board
- PS-6(3) - Access Agreements — Post-employment Requirements
- SR - Supply Chain Risk Management

C. Sensitivity Analysis

To help organizations quantitatively understand the content of each control pillar and maturity stage, the team performed text mining to extract information, including term frequency for each control pillar and the importance of each word in a control pillar or a maturity stage, relative to all NIST SP 800-53 controls.

1) Methodology: Term frequency examines the importance of words in a text by measuring how often certain words appear. The team first grouped the control text by their assigned control pillar (note that a control text can belong to one or multiple pillars). For each control pillar, the team treated it as a "bag of words", and after performing tokenization, removing stop words, and conducting lemmatization, the team found the respective term frequencies of the top 25 words for each control pillar.

Tokenization is the process of breaking a document into standardized word representations and splitting out separating punctuation. Each of these smaller units is called a token. Stop words are words that do not have actual meanings, like "the", "is", "in", "for", "where", "when", "to", etc. A word may exist in several inflected forms, and lemmatization removes inflectional endings only and returns the base or dictionary form of a word, known as the lemma. To perform tokenization, stop words removal, and lemmatization, the team utilized the NLTK platform to define functions that perform the steps. NLTK is a platform for building Python programs to work with human language data. It provides easy-to-use interfaces to over 50 corpora and lexical resources such as WordNet, and a suite of text processing libraries.

The team then utilized spaCy, an open-source software library for advanced natural language processing, to count the number of each "lemma", sort by frequency, and get the top 25 "lemma" for each control pillar.

While term frequency simply looks at the frequencies of words, TF-IDF (Term Frequency-Inverse Document Frequency) also looks at the "uniqueness" of a word to a control pillar or ZT maturity ranking. TF-IDF is the importance of a term is inversely related to its frequency across documents. TF (Term Frequency) is how often a term appears in a document, while IDF (Inverse Document Frequency) measures the relative rarity of a term in all control families/maturity rankings.

The higher the TF-IDF score, the more important a term is to a control family/maturity ranking, relative to other control families/maturity rankings. Hence, by examining the TF-IDF scores of words for each control family and maturity ranking, the team can understand how terms are more related to a specific group.

The team first defined two TF-IDF vectorizers and two corpora, one for control pillars and one for maturity rankings. After that, the team utilized the Scikit-learn library's fit_transform() function to perform fit and transform on the input data at a single time and convert the data points. To calculate the TF-IDF scores for each control pillar and maturity stage, the team utilized the NLTK platform to define functions that perform this step. The functions return the top 25 TF-IDF values in matrix rows, return them with their corresponding feature names, and get the top 25 TF-IDF features in a specific document (matrix row).

2) Findings and Analysis: Tables 1-6 represent the findings of the term frequency and TF-IDF analyses performed on the mapping.

TABLE I
TERM FREQUENCY BY CONTROL PILLAR: IDENTITY & DEVICE

Identity		**Device**	
Term	***Count***	***Term***	***Count***
assignment	334	assignment	438
system	260	system	383
information	114	component	121
access	99	following	105
security	83	information	91
individual	74	incident	78
policy	72	control	74
role	69	personnel	72
following	65	policy	70
personnel	63	security	69
frequency	59	frequency	63
organizational	51	role	62
component	50	access	56
privacy	48	organizational	56
user	45	response	56
assessment	43	service	52
procedure	41	plan	51
service	41	maintenance	51
contingency	40	change	48
selection	38	assessment	47
plan	35	procedure	46
account	34	selection	44
training	34	process	43
time	32	tool	43
audit	32	contingency	41

TABLE II
TERM FREQUENCY BY CONTROL PILLAR: NETWORK/ENVIRONMENT & APPLICATION WORKLOAD

Network / Environment		**Application Workload**	
Term	***Count***	***Term***	***Count***
assignment	613	system	450
system	472	assignment	361
component	163	security	130
information	155	component	122
following	149	information	94
control	114	service	88
security	104	following	85
policy	97	control	78
access	88	policy	70
personnel	85	privacy	64
frequency	83	frequency	53
incident	80	organizational	52
role	78	software	50
organizational	75	change	46
selection	74	role	46
response	61	assessment	46
time	56	personnel	45
change	55	developer	44
within	53	procedure	42
event	53	development	42
procedure	53	configuration	41
plan	53	design	41
mechanism	51	contingency	40
assessment	51	plan	39
external	50	process	38

TABLE III
TERM FREQUENCY BY CONTROL PILLAR: DATA

Data	
Term	***Count***
assignment	306
system	204
information	202
security	97
following	71
policy	70
control	63
frequency	52
privacy	49
processing	48
identifiable	48
access	46
personally	46
individual	45
assessment	42
contingency	42
role	39
component	37
medium	37
procedure	35
plan	34
training	34
organizational	34
selection	33
personnel	33

TABLE IV
TF-IDF BY CONTROL PILLAR (SCORE FROM HIGH TO LOW): IDENTITY, DEVICE, & NETWORK/ENVIRONMENT

Identity	**Device**	**Network / Environment**
assignment	assignment	assignment
system	system	system
information	following	information
access	informational	following
security	personnel	components
following	security	security
personnel	components	access
frequency	incident	personnel
policy	frequency	frequency
roles	organizational	organizational
individuals	policy	policy
organizational	roles	selection
privacy	access	systems
procedures	response	roles
contingency	component	incident
assessment	maintenance	controls
selection	procedures	response
controls	selection	time
training	contingency	procedures
audit	assessment	within
processing	plan	control
time	service	mechanisms
document	training	external
control	controls	physical
requirements	changes	monitoring

TABLE V
TF-IDF BY CONTROL PILLAR (SCORE FROM HIGH TO LOW): APPLICATION WORKLOAD & DATA

Application Workload	**Data**
system	assignment
assignment	information
security	system
information	security
following	following
component	policy
privacy	frequency
service	privacy
systems	processing
policy	identifiable
components	personally
frequency	access
organizational	contingency
software	systems
personnel	assessment
principle	roles
developer	individuals
roles	media
procedures	procedures
development	training
controls	organizational
design	personnel
assessment	components
contingency	selection
implement	control

TABLE VI
TF-IDF BY MATURITY RATING (SCORE FROM HIGH TO LOW)

Traditional	**Advanced**	**Optimal**
assignment	assignment	assignment
system	system	information
information	automated	security
security	information	system
following	access	privacy
frequency	security	lease
personnel	mechanisms	dynamic
organizational	components	attributes
access	follow	different
roles	time	domains
components	personnel	following
systems	actions	accounts
policy	control	devices
privacy	logon	mechanisms
procedures	changes	access
controls	tools	selection
selection	data	capability
physical	domains	filter
component	different	management
review	selection	identify
risk	incident	techniques
document	systems	accordance
service	processing	authentication
control	user	types
assessment	number	content

In general, there seem to be a few words that maintain greater association with each control pillar:

- Identity: access, roles, individuals
- Device: personnel, incident, maintenance
- Network/Environment: mechanisms, external, audit
- Application Workload: software, developer, firmware
- Data: identifiable, document, storage

Likewise, words like "automated" have a greater association with the "Advanced" maturity ranking, while "dynamic" and "analytics" are more related to the "Optimal" maturity ranking.

Future utilization of these methods suggests applying term frequency and TF-IDF findings to eliminate ambiguity in mapping and, consequently, ease the categorization of NIST SP 800-53 controls to ZT, based on the tendency of terms to coincide with different ZT control pillars and maturity rankings.

D. Web Application

The team built a web application to create an interactive way for organizations to search through the NIST SP 800-53 ZT mappings that were created in this research. This interactive method allows organizations to find relevant controls in a more digestible format. Companies can filter through the mappings depending on the CISA Zero Trust Maturity Model pillar and function of each NIST SP 800-53 control. The web application is a preliminary version of what the research team envisions becoming a fully fledged web application that contains logic to help entities implement ZT controls based on their organization's profile. The following sections walk through the functionality of this web application.

1) Methodology: To begin the web application, the data from the NIST SP 800-53 mapping needed to be loaded. The code shown in Listing 1 in the Appendix ensures that the NIST SP 800-53 mapping spreadsheet is converted into a dataframe that can be used to sort and filter through the rest of the web application. This code also loads in the primary and secondary functions that are used to filter the controls by their functionality. This allows organizations to filter the spreadsheet based on the functions of the controls.

The functionality of the code in Listing 2 in the Appendix is to add the CISA Zero Trust Maturity Model pillar. This functionality allows organizations to sort the NIST SP 800-53 ZT mapping based on the CISA Zero Trust Maturity Model pillars. This is essential for organizations as it allows them to identify the control to implement based off of the CISA Zero Trust Maturity Model pillars.

Listings 3 and 4 in the Appendix implement the filtering functionality throughout the web application. This code allows organizations to filter by Primary Function, Secondary Function, and CISA Zero Trust Maturity Model pillars. This is crucial for organizations to be able to granularly select ZT controls based on their organizations needs.

V. Project Limitations

The team initially identified the following deliverables in its Scope Statement to establish project scope and objectives:

1) *Control Decision Survey and Interviews*
   The project team plans to survey security managers across different companies and industries on their respective selection and implementation of controls (as they relate to ZT); the purpose of this is to understand how control selection criteria may differ by company and by industry.
2) *Decision Tree*
   The project team will conduct its own research of NIST SP 800-53, NIST SP 800-207, and other ZT publications, in addition to interviews with ZT SMEs (specific list of individuals to be interviewed is TBD). This analysis, taken in tandem with the above survey results, will be utilized to create a decision tree that can assist organizations in selecting the appropriate controls to implement given their respective business needs.
3) *Control Ranking Matrix*
   The project team will research different mathematical models and algorithms to use for control rankings, and utilize a selected algorithm for determining controls to implement against various security risks. This product is meant to quantitatively inform organizations on how to prioritize control selection and implementation.

Ultimately, within a few weeks of generating the preliminary Scope Statement, the team fine-tuned its initial deliverables to land on four key deliverables: (1) Zero Trust Survey, (2) NIST SP 800-53 / NIST SP 800-207 Control Mapping, (3) Sensitivity Analysis [based on the results of the mappings], (4) Zero Trust Decision Tree.

Throughout the continuation of the project, the team encountered several challenges, the first of which related to obtaining feedback for the distributed survey. Ultimately, the limitation here was the insufficient time to possess an adequate number of responses from security experts, related to their experience with and implementation of ZT within their organizations. Comprehensive

survey analysis assumes a much larger population of respondents is providing information, which would have been better attained through constant follow-up with subjects over the course of several months, as opposed to a few weeks.

The biggest limitation, however, was in the execution of the Control Mapping. There were a few factors that impacted this. The first was the breadth of NIST SP 800-53. With over 1,000 controls to analyze and map, the team found that it had not allotted enough time to sufficiently perform this task according to the initial scope and project schedule. Mapping efforts spanned several more weeks than had originally been anticipated, due to the extensiveness of the controls the team needed to address. Additionally, a vast majority of the NIST SP 800-53 controls do not promote either an advanced or optimal ZT maturity ranking, based on the definition of ZT maturity outlined by NIST SP 800-207; and, generally, the team discovered several discrepancies between NIST SP 800-53 and the expectation for ZT defined by CISA and the NIST SP 800-207 team.

Ultimately, the challenges the team faced in delivering the Control Mapping impacted its objective to develop an all-encompassing Decision Tree that could assist organizations in determining which controls to prioritize for migrating to ZT. These limitations altered the project scope, as the team saw a shift from aiming to create a tool to help organizations implement ZT, to discussing and presenting the areas in which NIST SP 800-53 and NIST SP 800-207 do not adequately overlap.

VI. Conclusion

The SEI capstone team drew many conclusions throughout the semester-long research on ZT. The biggest realization was that the industry lacks prescriptive guidance around ZT. The research presented in this report demonstrates that organizations are forward-thinking about ZTA, but generally lack guidance on how to achieve it. There is a large volume of information surrounding the ZT movement but limited information regarding its actual implementation. Tooling is needed for organizations to implement ZT into their environments. The CMU SEI capstone research team has diligently mapped all NIST SP 800-53 controls to the CISA Zero Trust Maturity Model to further this goal. The CMU research team believes that the research and documentation provided in this report can serve as a foundation for organizations to select and implement controls based on their unique ZT and broader organizational needs. This report can be expanded upon to develop a framework to implement ZT in a wholistic and quantitative manner, with customized guidance for individual organizations.

A. Future Work

In accordance with the team's initial scope for the project, it proposes continued revision of its mapping efforts and interactive web application, as well as the development of the following tools to further advance ZTA standardization across organizations and industries:

1) Decision Model: The SEI capstone team has determined that the most helpful tool for organizations would be a decision model based on the mapping provided in this research. This decision model would intake a company's profile and determine its maturity rating according to the CISA Zero Trust Maturity Model. However, it would not just rate each company's overall maturity but also their maturity according to each of the ZT pillars, such as Identity and Application Workload. The tool would also assess the company's market segment, compliance requirements, and technological assets to prioritize ZT controls by ZT pillar. The outcome would be a prioritized list of controls by Zero Trust pillar that the company should implement. This decision model would help a company quantify its implementation roadmap and give them prescriptive guidance based on the companies' specific needs.

2) ROI Model: Furthermore, some companies' security strategy is focused on the return on investment (ROI) of controls. During the research, the CMU SEI Capstone team discovered that the thought leaders on ZT believe that a method of calculating the ROI of ZT controls would significantly improve ZT adoption throughout the industry. Essentially, this tool would allow security leaders to forecast the cost of implementing a ZT control and see their expected ROI. This tool, just like the decision model, would intake a company's profile to make the recommendation specific for that organization. The ROI model would help security leaders prioritize their implementation of ZT controls based on their business's needs.

APPENDIX

```
df = pd.read_excel("map.xlsx")
column_name_primary_f =
    'Primary Function'
column_name_secondary_f =
    'Secondary Function'
primary_functions =
    df[column_name_primary_f
    ].drop_duplicates(
        ).values.tolist()
secondary_functions =
    df[column_name_secondary_f
    ].drop_duplicates(
        ).values.tolist()
```

Listing 1: Variables and document calling.

```
st.sidebar.write("Control Family")

control_family_names = ["Identity",
    "Device","Network/Environment",
    "Application Workload", "Data"]
controlFamiliesSelections = [
    st.sidebar.checkbox(s, value=True
    ) for s in control_family_names
]
controlFamiliesSelections = [
    1 if i else 0 for i
    in controlFamiliesSelections]
[id, dev, netenv, appwl, dt] =
    controlFamiliesSelections
```

Listing 2: Control pillar names.

```
def controlFilter(df):
    indexes = [df[
        control_family_names[i]] ==
        controlFamiliesSelections[i]
        if
        controlFamiliesSelections[i]
        else
        df[control_family_names[i]] ==
        df[control_family_names[i]]
    for i in range(len(
        controlFamiliesSelections))
    ]
    return df.loc[
        reduce(lambda x, y : x & y,
        indexes)
     ]
st.sidebar.write("---------------")

pf_enable= st.sidebar.checkbox(
    "Filter by Primary Function",
    value=False)
if pf_enable:
    pf_select = st.sidebar.selectbox(
        "Primary Function",
        primary_functions)
sf_enable = st.sidebar.checkbox(
    "Filter by Secondary Function",
    value=False)
if sf_enable:
    sf_select = st.sidebar.selectbox(
        "Secondary Function",
        secondary_functions)

def functionFilter (df):
    retDf = df
    if pf_enable:
        retDf = retDf[retDf[
            column_name_primary_f] ==
                pf_select]
    if sf_enable:
        retDf = retDf[retDf[
            column_name_secondary_f
            ] == sf_select]
    return retDf
```

Listing 3: Filter function, part 1.

```
print(controlFamiliesSelections)
finalDF = functionFilter(
    controlFilter(df))
st.write("Total {} lines".format(
    len(finalDF)))
st.dataframe(finalDF, width=800,
    height=600)
# st.table(finalDF)

st.write("\n\n".join(finalDF[
    "Control Text"].values.tolist()))
```

Listing 4: Filter function, part 2.

ACKNOWLEDGMENT

The SEI capstone team would like to express its deepest appreciation to its project sponsor the Software Engineering Institute, as well as its client Mr. Brett Tucker, who provided the opportunity to contribute to current ZT research. Its gratitude also extends to Professor Randy Trzeciak, who offered advice and resources throughout the semester.

Thank you to all of the faculty and staff at Heinz College, who provided support and guidance to each of the team members during the course of their graduate studies.

REFERENCES

[1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, NIST Special Publication 800-207: Zero Trust Architecture. Stafford, VA: National Institute of Standards and Technology, 2020. doi: 10.6028/nist.sp.800-207.

[2] K. Raina, "Zero Trust Security Explained — Principles of the Zero Trust Model," Crowdstrike, May 06, 2021. https://www.crowdstrike.com/cybersecurity-101/zero-trust-security/.

[3] IBM, "Cost of a Data Breach Study," Ibm.com, 2021. https://www.ibm.com/security/data-breach.

[4] L. M. Pelzer, "The True Cost of Cybersecurity Incidents: The Problem," Palo Alto Networks, Jun. 25, 2021. https://www.paloaltonetworks.com/blog/2021/06/the-cost-of-cybersecurity-incidents-the-problem/.

[5] M. Kaufmann, "How Privileged Access Works with Zero Trust," Saviynt, Dec. 03, 2020. https://saviynt.com/how-privileged-access-works-with-zero-trust/.

[6] National Security Agency, "Embracing a Zero Trust Security Model," U.S. Department of Defense, Feb. 2021. [Online]. Available: https://media.defense.gov/2021/Feb/25/2002588479/-1/-1/0/CSI_EMBRACING_ZT_SECURITY_MODEL_UOO115131-21.PDF.