

Understanding Zero Trust Security Implementations via the MITRE ATT&CK and D3FEND Frameworks: Uncovering Trends Across a Decade of Breaches

Philip Menard
The University of Texas at
San Antonio
philip.menard@utsa.edu

Elizabeth Reyes
The University of Texas at
San Antonio
elizabeth.reyes4@my.utsa.edu

Raymond M. Bateman
U.S. Army Combat Capabilities
Development Command
Raymond.M.Bateman4.civ@army.mil

Abstract

Information sharing is paramount to operating within the modern business domain. However, with information sharing comes the risk of data breaches. One of the key challenges facing organizations is the ability to trace, and therefore trust, digital information flows. Due to its central philosophy of verifying network traffic before trusting it, zero trust security is an approach to cyber defense architecture that is rapidly gaining popularity across organizations. Although fully adopting zero trust should greatly reduce an organization's likelihood of suffering a breach, organizations adopt zero trust in varying degrees. In this manuscript, we aim to better understand how zero trust has been adopted over the last decade, using Verizon's Data Breach Incident Report dataset as a representative sample whereby we may infer lack of zero trust adoption via observable breaches. We find that certain aspects are positively correlated with breach occurrences, while others are negatively associated.

1. Introduction

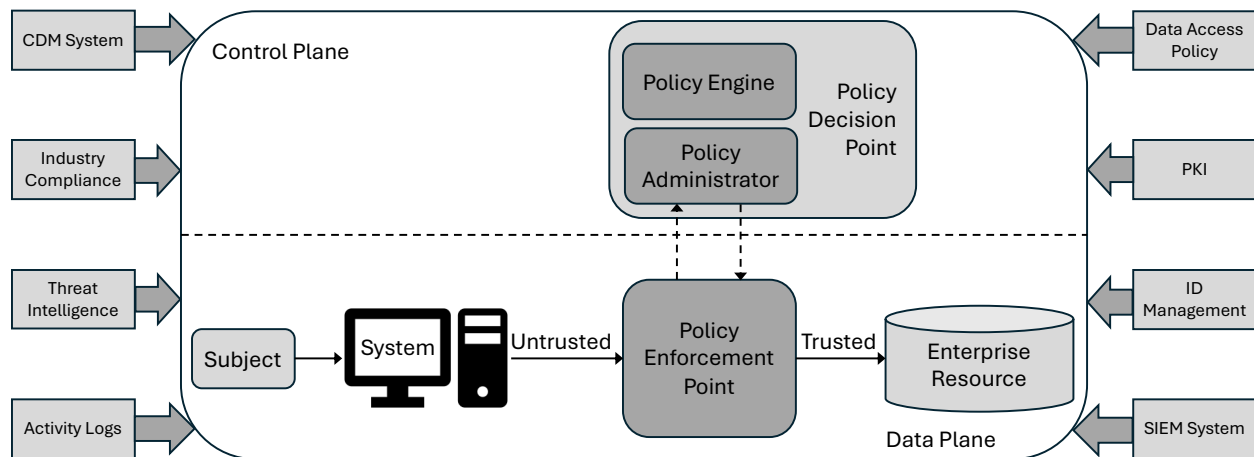
Securing an enterprise is a challenging task. This difficulty stems from various factors, including the intricate nature of IT and application infrastructures, the vast and rapid user access, and the inherently adversarial landscape of information security (Chen et al., 2019; Compastie et al., 2016). Additionally, the openness of most enterprise networks exacerbates these challenges. Without implementing the principle of least privilege at both network and application levels, organizations expose themselves to significant vulnerabilities (Mehraj & Bandy, 2020; Zaheer et al., 2019). Traditional security and networking systems, still widely used, perpetuate this flawed model.

Zero Trust security revolutionizes this approach by adopting a modern security paradigm that rigorously

applies the principle of least privilege to networks and applications (Kerman et al., 2020). Under this model, unauthorized users and systems are completely denied access to enterprise resources, while authorized users are granted only the minimum level of access required (Campbell, 2020). This approach ensures that enterprises become safer, more secure, and more resilient (Cunningham, 2020). Moreover, Zero Trust enhances efficiency and effectiveness by automating the enforcement of dynamic and identity-centric access policies (Adahman et al., 2022).

Due to their flexible deployment and scalability, cloud enterprise architectures are suitable for zero trust implementation (Albuali et al., 2020; Mehraj & Bandy, 2020). Security administrators can easily adjust networking architectures and deploy additional countermeasures as needed, while ensuring zero trust compliance through their cloud service provider contracts. However, most organizations do not rely solely on cloud environments for their information technology infrastructure, but rather a hybrid approach that blends on-premises legacy systems with newer cloud-based architectures (Balaouras et al., 2018; Rose et al., 2020). While cloud environments offer abstracted logical architectures that encourage adopting newer security approaches as they emerge, such as zero trust (Yan & Wang, 2020), on-premises systems require significant resources for security patching, software maintenance, security testing, and ensuring regulation compliance (Bertino & Brancik, 2021). These activities may also result in sub-optimal retrofitting that may not be fully compliant with modern approaches like zero trust (Bertino & Brancik, 2021; Rose et al., 2020). As a result, organizations who adopt a hybrid cloud/on-premises architecture may inherently be adopting a weaker cybersecurity posture, leaving them vulnerable to exploits that would otherwise be more difficult to execute in environments that fully adopt zero trust principles.

This transition from traditional on-premises approaches to zero trust architecture (or some



CDM = continuous diagnostics and mitigation; PKI = public key infrastructure; ID = identification;
 SIEM = security information and event management

Figure 1. Core zero trust logical components (adapted from Rose et al., 2020)

combination thereof) can potentially be tracked through the incidents that have been observed and reported since the initial publication of the zero trust guiding principles. Verizon’s Data Breach Incident Report (DBIR) is an annual assessment of the cyber threat landscape facing organizations (Verizon Enterprise Solutions, 2024). The data that informs the DBIR is updated and made public by Verizon via GitHub, allowing cybersecurity analysts and researchers to glean critical findings on understanding attack motivations and recommending better cybersecurity best practices (Verizon, 2024). We posit that this dataset may be useful in determining how well zero trust principles have addressed the occurrence of cybersecurity incidents and pose the following research question:

RQ: How well do zero trust principles explain the occurrences of data breach incidents?

In this study we use the DBIR dataset to examine zero trust adoption over the last decade. Our research offers a novel mapping of zero trust principles to the DBIR dataset, allowing us to use successful attack types to infer zero trust adoption failure. Further, we use our analysis to offer recommendations for current hybrid cloud/on-premises approaches, highlighting their strengths and weaknesses within the context of cybersecurity, using the MITRE ATT&CK and D3FEND frameworks as guides.

2. Literature Review

2.1. Zero Trust Security

In 2010, Forrester introduced the concept of “Zero Trust” in a whitepaper that encapsulated ideas circulating in the industry for some time (Kindervag et

al., 2010). In the initial development and implementation of networking topologies, especially those that pre-dated the widespread adoption of the Internet, internal traffic was assumed to be trustworthy, and there was little to no opportunity for unverified traffic to be allowed into the network. Security measures, if implemented at the network level at all, were relegated to the edge of the network with little to no internal verification mechanisms. As organizations expanded their operations, and resulting information processing, beyond the confines of local networks, there was a pertinent need to secure internal resources from outside untrusted sources; relatedly, internal network traffic could no longer be assumed to be inherently trustworthy.

Forrester’s approach outlined a departure from the concept of a rigid perimeter, advocating instead for a methodology that necessitated the inspection and understanding of network elements before granting trust and access. Over time, Forrester refined this concept into what is now known as the Zero Trust eXtended (ZTX) Framework, with Data, Workloads, and Identity identified as core components of Zero Trust (Cunningham, 2018). Around the same time as Forrester’s 2010 publication of Zero Trust principles, Google launched its internal BeyondCorp initiative, which implemented a form of Zero Trust and eliminated the traditional enterprise network boundary (Ward & Beyer, 2014). Concurrently, in 2014, the Cloud Security Alliance introduced the Software Defined Perimeter architecture, offering a concrete specification for a security system supporting Zero Trust principles (Samani et al., 2014).

In the subsequent years following Forrester’s and Google’s advancements, the emphasis on Zero Trust persisted across the industry, with the US National

Institute of Standards and Technology (NIST) releasing a Zero Trust Architecture publication (Rose et al., 2020) and launching an associated US National Cybersecurity Center of Excellence project in 2020. Broadly speaking, zero trust principles are enforced through the implementation of a policy enforcement point (PEP) and a policy decision point (PDP). All system traffic is considered untrusted until verified by the PEP and PDP. The PEP is responsible for accepting access requests from unverified sources. The PEP then checks the PDP to determine whether the as-yet-unverified source should be trusted to access the requested resource. Zero trust processes are also logically split into the data plane and the control plane. Untrusted sources, the PEP, and secured resources are located on the data plane of the zero trust infrastructure, while the PDP and its associated processes are relegated to the control plane. This basic logic is extensible to incorporate a variety of common enterprise and/or security systems, including ID management, SIEM, PKI, threat intelligence, and more. Garbis & Chapman (2021) provide a comprehensive guide for various zero trust implementation strategies. The specific components of NIST's zero trust framework are depicted in Figure 1.

2.2. Recent Literature on Zero Trust

While the NIST publication has helped unify varied zero trust approaches under common components and nomenclature, many differences in implementation still exist in practice. These differences have opened interesting avenues for research within the context of zero trust implementation. Since the initial publication of the overarching zero trust principles, researchers have examined zero trust approaches in several ways. Much of the work in this area has been focused on technical specifications and examining the boundary conditions of zero trust effectiveness. Studies have determined that zero trust is a dynamic architectural model (Yeoh et al., 2023) that largely relies on micro-segmentation (Rose et al., 2020) and a thorough application of the principle of least privilege (Campbell, 2020; Yan & Wang, 2020) and system logging and analysis (Yan & Wang, 2020) to reach its full potential of comprehensive cybersecurity coverage. Recent works have also identified the factors that have the largest impact on successful zero trust implementations, equipping organizations with tools to assess their current security posture (Yeoh et al., 2023) and highlight areas of further inquiry (Buck et al., 2021).

Researchers have identified the analysis of zero trust shortcomings as critical to advancing knowledge in this area (Buck et al., 2021). Despite the recent

wealth of research regarding the conceptual and technical deployments of zero trust, researchers have yet to examine zero trust from the vantage of observed attacks, which may offer important insights on how effectively zero trust principles have been adopted by organizations.

3. Methods

To examine the relationship between zero trust principles and observed data breach incidents, we utilized a combination of the VERIS Community Database and the MITRE ATT&CK and D3FEND frameworks, each described in further detail below.

3.1. Verizon DBIR and VERIS

For the last several years, Verizon has compiled data breach information, summarizing the year in cyber threats in its Data Breach Incident Report (DBIR). The DBIR provides an annual snapshot of the most pertinent cybersecurity threats posed to organizations, with useful summaries of the most common attack types, breakdowns of attacks by industry sector, and overarching trends in the threat landscape.

The vocabulary for event recording and incident sharing (VERIS) was developed by Verizon as a way to standardize information sharing when reporting cybersecurity incidents. They also use this framework as the basis for compiling the incident data that informs their annual DBIR. Verizon has made their incident database, the VERIS Community Database (VCDB), public through GitHub (<https://github.com/vz-risk/VCDB>), with the VCDB's incidents coded using attack capability identifiers defined within the VERIS framework. Because it is based on the extensive research that informs the annual DBIR, the VCDB has emerged as one of the most comprehensive and publicly available dataset related to data breaches, thus providing a representative sample of cyber attacks from which to glean overall trends and insights. The VCDB contains information on 10,262 breach incidents, with nearly all incidents in the dataset dating from 2014 to the date that we downloaded the database for analysis (April 21, 2024). Each incident may be classified with multiple VERIS capability identifiers; indeed, no entry in the VCDB consists of just one attack capability. The company who suffered the data breach is anonymized in the dataset, but categorical identifiers (such as industry sector) are included. Analysis related to company type is outside the scope of this research but will be utilized in future research.

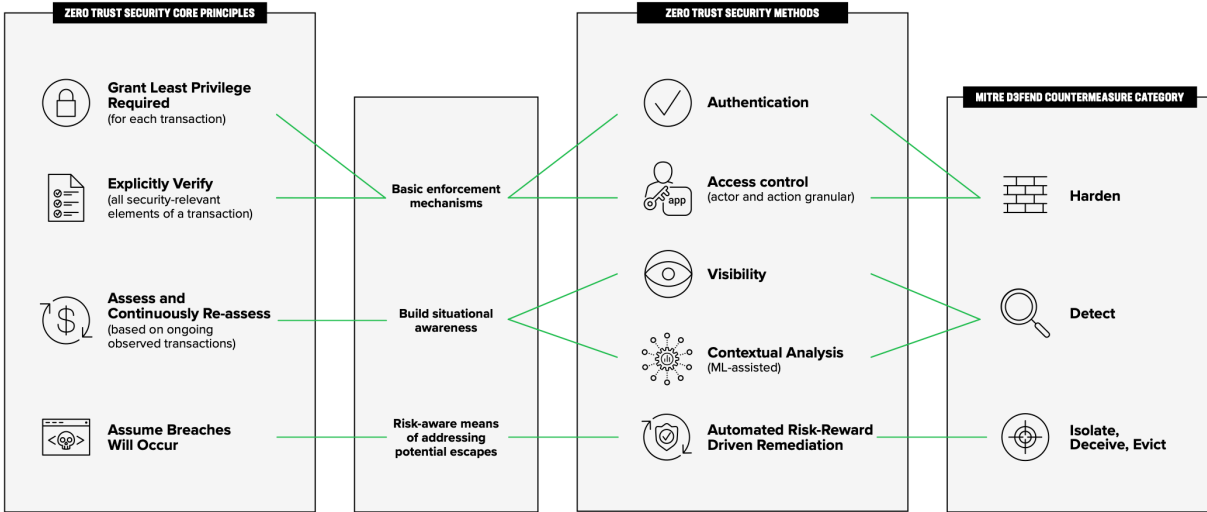


Figure 2. Mapping between zero trust security techniques and MITRE D3FEND countermeasure categories (Source: Tyagi & Arora, 2022)

3.2. MITRE ATT&CK

Table 1. MITRE ATT&CK tactic definitions

Tactic	Definition
Reconnaissance	The adversary is trying to gather information they can use to plan future operations
Resource Development	The adversary is trying to establish resources they can use to support operations
Initial Access	The adversary is trying to get into your network
Execution	The adversary is trying to run malicious code
Persistence	The adversary is trying to maintain their foothold
Privilege Escalation	The adversary is trying to gain higher-level permissions
Defense Evasion	The adversary is trying to avoid being detected
Credential Access	The adversary is trying to steal account names and passwords
Discovery	The adversary is trying to figure out your environment
Lateral Movement	The adversary is trying to move through your environment
Collection	The adversary is trying to gather data of interest to their goal
Command and Control	The adversary is trying to communicate with compromised systems to control them
Exfiltration	The adversary is trying to steal data
Impact	The adversary is trying to manipulate, interrupt, or destroy your systems and data

Similar to VERIS, the MITRE Corporation has developed its own knowledge base, ATT&CK, for cataloging cyber adversary behavior (MITRE Corporation, 2024a). ATT&CK classifies attacks using broader “tactic” categories, under which a variety of techniques are housed, with some techniques having specific sub-techniques cataloged as well. The ATT&CK tactic categories are

Reconnaissance, Resource Development, Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, Collection, Command and Control, Exfiltration, and Impact. There are 527 unique ATT&CK techniques and sub-techniques. While most techniques and sub-techniques belong to just one tactic category, several techniques are mapped to multiple tactical types. Using this design, MITRE developed ATT&CK to analyze cybersecurity threats in both a high-level manner and in finer-grained detail. The definitions of MITRE’s ATT&CK tactic categories are provided in Table 1.

3.3. MITRE D3FEND

Complementing MITRE ATT&CK is MITRE’s catalog of applicable cybersecurity countermeasures, called D3FEND (MITRE Corporation, 2024b). The MITRE D3FEND framework is organized similarly to MITRE ATT&CK, whereby grand-level countermeasure “tactics” house specific countermeasure “techniques.” The D3FEND tactic categories are Model, Harden, Detect, Isolate, Deceive, Evict, and Restore. There are 119 unique D3FEND techniques in the latest version of the framework.

Additionally, MITRE D3FEND plots the relationships between D3FEND and ATT&CK techniques. These relationships are depicted graphically on the D3FEND website and are available as downloadable JSON. A single D3FEND technique may be mapped to multiple ATT&CK techniques, and vice versa. As of the writing of the manuscript, MITRE has only mapped D3FEND techniques to

enterprise ATT&CK techniques; however, this is not a significantly limiting factor, as the majority of incidents in the VCDB involve enterprise techniques (9,477 entries out of 10,262 total entries). The definitions of MITRE’s D3FEND tactic categories is provided in Table 2.

Table 2. MITRE D3FEND tactic definitions

<i>Tactic</i>	<i>Definition</i>
Model	used to apply security engineering, vulnerability, threat, and risk analyses to digital systems; accomplished by creating and maintaining common understanding of the systems being defended, operations on those systems, actors using the systems, and relationships and interactions between these elements
Harden	used to increase the opportunity cost of computer network exploitation; differs from Detection in that it generally is conducted before a system is online and operational
Detect	used to identify adversary access to or unauthorized activity on computer networks
Isolate	creates logical or physical barriers in a system which reduces opportunities for adversaries to create further accesses
Deceive	used to advertise, entice, and allow potential attackers access to an observed or controlled environment
Evict	used to remove an adversary from a computer network
Restore	used to return the system to a better state

Finally, Tyagi and Arora (2022) used the MITRE D3FEND tactics to create a mapping of countermeasures to zero trust principles (see Figure 2). The authors identified the following MITRE D3FEND tactics as the most relevant to zero trust: Harden (associated with the zero trust principles “Grant least privilege required” and “Explicitly verify”), Detect (associated with the zero trust principle “Assess and continuously re-assess”), Isolate, Deceive, and Evict (each associated with the zero trust principle “Assume breaches will occur”). Because VERIS and MITRE D3FEND are both mapped to MITRE ATT&CK, Tyagi and Arora’s zero trust mapping provides the final link needed to connect zero trust principles to the VCDB.

4. Data Analysis and Results

We began by compiling a novel dataset that incorporates MITRE ATT&CK and D3FEND classifiers for each of the reported incidents in the VCDB. Using the most recent VCDB dataset as a baseline (downloaded on April 21, 2024), we added variables to the dataset representing whether a specific MITRE ATT&CK or D3FEND technique would be related to the VERIS capability. Using R Studio as our statistical software (RStudio Team, 2020), we

compiled our expanded VCDB by determining which VERIS capability identifiers were associated with a specific incident and then marking the MITRE ATT&CK techniques that were associated with the VERIS capability. The R script can be found at https://github.com/menap553/VERIStoMITRE_HICSS58.

4.1. Data Transformation

After linking the various mappings together and creating new variables within the base VCDB dataset, we mutated the data such that we could track the count and frequency of each MITRE ATT&CK technique, while also tracking which D3FEND countermeasures corresponded with the specific attack technique. We plotted these occurrences using histograms, determining that the attacks were clustered into three distinct frequency groupings. We found that a subset of 62 ATT&CK techniques appeared in the VCDB in exceedingly high rates, thus representing a high-occurrence outlier subset (7,100+ appearances in the 9,477 VCDB incidents). Because of the small sample size of this group, these incidents were excluded from the remainder of the analysis. The remaining ATT&CK techniques were separated into “low-rate” (ranging 1-500 appearances) and “medium-rate” (ranging 900-2,800 appearances) subsets, resulting in the low-rate subset containing 200 ATT&CK techniques and the medium-rate subset containing 92 techniques. The ranges included in these rate levels are not continuous because no attacks occurred at those frequencies (i.e., there were no attacks that occurred between 500 and 900 times, no between 2,800 and 7,100 times).

4.2. Negative Binomial Regression Results

To analyze the association between zero trust-oriented MITRE D3FEND tactics and the number of occurrences of MITRE ATT&CK techniques, we used negative binomial regression. This statistical technique is appropriate for models where the dependent variable is a count of occurrences (Gardner et al., 1995); because the dependent variable for each of our models is the number of times an ATT&CK technique appeared in the VCDB. In each regression model, we included all D3FEND tactic categories as independent variables. We also include each ATT&CK tactic category to control for whether techniques from a specific category were more likely to occur.

For the low-rate subset, ATT&CK techniques associated with MITRE D3FEND countermeasures in the Deceive category were more likely to occur (est. = 0.523; p = 0.030), whereas ATT&CK techniques

associated with MITRE D3FEND countermeasures in the Evict category were less likely to occur (est. = -0.478; p = 0.040). Whether an ATT&CK type was associated with D3FEND countermeasures in the Harden, Detect, or Isolate categories did not influence the ATT&CK technique's occurrence rate.

Outside of zero trust-oriented factors, an attack technique's association with the D3FEND Model category negatively influenced number of occurrences (est. = -0.683; p < 0.001), while the Restore category positively influence attack technique occurrence (est. = 0.630; p = 0.018). An attack technique's classification as Command and Control (est. = 0.980; p = 0.027), Credential Access (est. = 1.187; p < 0.001), Defense Evasion (est. = 0.438; p = 0.009), Impact (est. = 1.874; p < 0.001), Lateral Movement (est. = 1.767; p < 0.001), and Privilege Escalation (est. = 1.171; p < 0.001) positively affected attack occurrence. Table 3 provides details on the negative binominal regression analyses for the low-rate subset (D3FEND tactics related to zero trust are highlighted).

Table 3. Negative binomial regression effects on low-rate MITRE ATT&CK techniques in the VCDB

	Estimate	S.E.	Z value	p value
(Intercept)	3.640	0.260	14.019	< 0.001
D3FEND Tactic Categories				
Model	-0.683	0.202	-3.377	< 0.001
Harden	-0.202	0.223	-0.908	0.364
Detect	0.171	0.229	0.748	0.454
Isolate	-0.080	0.162	-0.493	0.622
Deceive	0.523	0.241	2.171	0.030
Evict	-0.478	0.232	-2.057	0.040
Restore	0.630	0.267	2.361	0.018
ATT&CK Tactic Categories				
Collection	0.366	0.314	1.165	0.244
Command and Control	0.980	0.444	2.206	0.027
Credential Access	1.187	0.248	4.795	< 0.001
Defense Evasion	0.438	0.168	2.607	0.009
Discovery	0.193	0.267	0.722	0.470
Execution	0.382	0.289	1.323	0.186
Exfiltration	---	---	---	---
Impact	1.874	0.389	4.814	< 0.001
Initial Access	0.463	0.283	1.636	0.102
Lateral Movement	1.767	0.279	6.334	< 0.001
Persistence	0.236	0.195	1.209	0.227
Privilege Escalation	1.171	0.183	6.415	< 0.001
Reconnaissance	---	---	---	---
Resource Development	---	---	---	---

For the medium-rate subset, ATT&CK techniques associated with MITRE D3FEND countermeasures in the Deceive category were again more likely to occur (est. = 0.242; p = 0.003), whereas ATT&CK

techniques associated with MITRE D3FEND countermeasures in the Harden (est. = -0.295; p < 0.001) and Isolate categories (est. = -0.172; p < 0.001) were less likely to occur. Whether an ATT&CK type was associated with D3FEND countermeasures in the Detect or Evict categories did not influence the ATT&CK technique's occurrence rate. Table 4 provides details on the negative binominal regression analyses for the medium-rate subset.

Besides zero trust-oriented factors, an attack technique's association with the D3FEND Model category negatively affected number of occurrences (est. = -0.104; p = 0.038), while the Restore category positively affected attack technique occurrence (est. = 0.272; p = 0.019). An attack technique's classification as Defense Evasion (est. = 0.264; p < 0.001), Impact (est. = 0.568; p < 0.001), and Privilege Escalation (est. = 0.228; p < 0.001) were positively associated with attack occurrence. Command and Control (est. = -0.217; p = 0.004) and Persistence (est. = -0.182; p = 0.002) were negatively associated with attack occurrence. Table 4 provides details on the negative binominal regression analyses for the medium-rate subset (D3FEND tactics related to zero trust are highlighted).

Table 4. Negative binomial regression effects on medium-rate MITRE ATT&CK techniques in the VCDB

	Estimate	S.E.	Z value	p value
(Intercept)	3.640	0.260	14.019	< 0.001
D3FEND Tactic Categories				
Model	-0.104	0.050	-2.072	0.038
Harden	-0.295	0.082	-3.603	< 0.001
Detect	0.004	0.085	0.049	0.961
Isolate	-0.172	0.052	-3.293	< 0.001
Deceive	0.242	0.081	2.984	0.003
Evict	-0.145	0.097	-1.503	0.133
Restore	0.272	0.116	2.351	0.019
ATT&CK Tactic Categories				
Collection	---	---	---	---
Command and Control	-0.217	0.076	-2.854	0.004
Credential Access	0.035	0.115	0.306	0.760
Defense Evasion	0.264	0.051	5.202	< 0.001
Discovery	---	---	---	---
Execution	-0.074	0.081	-0.911	0.362
Exfiltration	---	---	---	---
Impact	0.568	0.095	6.01	< 0.001
Initial Access	-0.044	0.070	-0.627	0.530
Lateral Movement	0.080	0.125	0.64	0.522
Persistence	-0.182	0.059	-3.052	0.002
Privilege Escalation	0.228	0.051	4.48	< 0.001
Reconnaissance	---	---	---	---
Resource Development	---	---	---	---

5. Discussion

5.1 Overall Findings

In this study, we aimed to answer the following question: *How well do zero trust principles explain the occurrences of data breach incidents?* Our results, based on the past ten years of VCDB breach reporting, is that zero trust's ability to explain attack occurrences is mixed. Some principles are correlated with low-rate attacks, while others are associated with medium-rate attacks. In each frequency grouping, certain ATT&CK tactics also demonstrate significant influence on rates of specific ATT&CK techniques.

Our analyses of the low-rate and medium-rate ATT&CK subsets yielded interesting findings. D3FEND techniques classified under the Harden tactic category are used to increase the opportunity cost of computer network exploitation. D3FEND techniques classified under the Isolate tactic category create logical or physical barriers in a system which reduces opportunities for adversaries to create further accesses. Most attacks would need to circumvent host hardening techniques and be able to move throughout an organization's infrastructure to successfully perpetrate its intended action. Thus, observing a negative association between both Harden and Isolate techniques and attack occurrences in the medium-rate subset is not surprising. In the low-rate subset, both of these D3FEND tactic categories were negatively related to attack occurrences but not to a significant degree. Based on these findings, organizations generally seem to be implementing countermeasures associated with host hardening and adversary isolation but should also ensure that they are protected against less-common attack types.

D3FEND techniques classified under the Evict tactic category remove an adversary from a computer network. Techniques under this category are negatively associated with attack occurrences for each of the subsets, but only to a significant degree for the low-rate subset. This finding may offer a partial explanation as to why the attack techniques in the low-rate subset are less common – an attack cannot be successful if the adversary is at risk of being removed from the network. Thus, attack techniques which can effectively evade eviction would occur at higher rates. Still, the negative correlation between Evict techniques and attack occurrences offer some evidence that organizations are implementing such countermeasures.

D3FEND techniques classified under the Deceive tactic category are used to advertise, entice, and allow potential attackers access to an observed or controlled

environment. Intuitively, such techniques are indeed positively and significantly associated with attack occurrences in both the low-rate and medium-rate attack subsets. This finding demonstrates that organizations are effectively attracting attacker activity that may be leveraged for subsequent intelligence gathering. However, because the attacks records in the VCDB are related to successful breaches, this finding may also indicate that organizations are properly managing their deception techniques, such that adversaries are still able to perpetrate attacks despite the deception.

D3FEND techniques classified under the Detect tactic category are used to identify adversary access to or unauthorized activity on computer networks. In each of the subset, the Detect category exhibited a positive, non-significant correlation with attack occurrences. Whether an attack is associated with a detection-oriented countermeasure does not seem to influence its rate of occurrence. This finding highlights an important area of improvement for organizations.

The two D3FEND tactic categories not linked to zero trust principles, Model and Restore, were each significantly related to attack occurrences. D3FEND techniques classified under the Model tactic category are used to apply security engineering, vulnerability, threat, and risk analyses to digital systems. A common refrain among cybersecurity professionals is the emphasis on strategic planning. Our finding that techniques under the Model category are negatively associated with attack occurrences in both subsets is not surprising. Conversely, D3FEND techniques classified under the Restore tactic category are used to return the system to a better state. By definition, these techniques are used to remediate successful attacks, so the significant positive correlation between the Restore category and attack occurrences is also intuitive.

Several ATT&CK tactic categories were correlated with increased attack technique occurrences. The Defense Evasion, Impact, and Privilege Escalation categories were each significantly and positively associated with attack technique occurrences for both the low-rate and medium-rate subsets. Because the Impact and Privilege Escalation categories represent common overarching goals for most cyber adversaries, their significant correlations with attack occurrences is expected. Because Defense Evasion is directly related to techniques designed to avoid detection, this finding provides further evidence of organizations' lack of adequate detection countermeasures.

Other ATT&CK tactic categories demonstrated mixed results. Credential Access and Lateral

Movement were only significantly correlated with attack technique occurrences in the low-rate subset. Techniques under the Credential Access category represent attacks where adversaries are trying to steal account names and passwords, while techniques under the Lateral Movement category represent attacks where adversaries are trying to move through an organization's environment. Interpreted together, these findings may indicate that attacks that comprehensively compromise an organization's infrastructure are rarer, while most attacks seem to be more targeted and short-term in nature. Command and Control was positively associated with attack technique occurrences in the low-rate subset but was negatively correlated with occurrences in the medium-rate subset. Similarly, Persistence was negatively related to attack technique occurrences in the medium-rate subset. Techniques under these categories may represent more advanced breach techniques and are therefore rarer than techniques that are more easily executable or require less adversarial effort.

5.2 Implications on Research

Altogether, the above findings illustrate interesting patterns of zero trust adoption based on attack occurrences over the last decade. The Harden tactic category from the MITRE D3FEND framework is associated with the "Grant least privilege required" and "Explicitly verify" zero trust principles. Based on our findings, organizations seem to be doing a reasonable job adopting these principles. The Isolate, Deceive, and Evict categories are each associated with the "Assume breaches will occur" zero trust principle. The Isolate and Evict categories appear to complement one another; each are negatively associated with attack occurrences, while the associations are significant in opposite subsets (Evict for the low-rate subset; Isolate for the medium-rate subset). However, the effectiveness of countermeasures under the Deceive category require further investigation, as it is not currently clear whether its positive effect on attack occurrences is attributable to adequate deceptive countermeasures or simply the organization's inherent attractiveness as a breach target. Finally, the Detect tactic category is associated with the "assess and continually re-assess" zero trust principle. Organizations seem to have given countermeasures in this category the least amount of attention, as evidenced by the lack of significant effects observed in either data subset. Thus, the implementation of detection-oriented countermeasures appears to be the most obvious area of improvement for organizations' security profiles.

5.3 Limitations and Future Research

A limitation of our work is a reliance on analyzing trends associated with breaches that have been observed, which inherently excludes breaches that have gone undetected. Future research may supplement our findings with additional intelligence sources, such as hacker forums, which may discuss successful breaches that have remained untraceable. Although the level of detail that may be gleaned from such sources may not be as comprehensive as what is provided in the VCDB, the information may still prove valuable in fully understanding the data breach landscape. Another limitation is our exclusive use of the VCDB as our data source for breach information. Future researchers may want to replicate our approach using varying breach datasets to determine if the findings presented here are still applicable to other breach incidents.

6. Conclusion

Due to its emphasis on robust identity-based trust, zero trust security has the potential to substantially strengthen organizations' cybersecurity postures. However, adoption of zero trust has thus far been varied and often implemented in a piecemeal fashion. By linking zero trust principles and associated countermeasures with observed cyber attack techniques, our research sheds light on important areas of improvement for organizational cybersecurity.

7. References

- Adahman, Z., Malik, A. W., & Anwar, Z. (2022). An analysis of zero-trust architecture and its cost-effectiveness for organizational security. *Computers & Security, 122*, 102911.
- Albuali, A., Mengistu, T., & Che, D. (2020). ZTIMM: A zero-trust-based identity management model for volunteer cloud computing. *Cloud Computing—CLOUD 2020: 13th International Conference, Held as Part of the Services Conference Federation, SCF 2020, Honolulu, HI, USA, September 18-20, 2020, Proceedings 13*, 287–294.
- Balaouras, S., Cunningham, C., & Cerrato, P. (2018). *Five Steps to a Zero Trust Network: Road Map: The Security Architecture and Operations Playbook*. Forrester Research. <https://www.forrester.com/report/Five+Steps+To+A+Zero+Trust+Network/-/E-RES120510>
- Bertino, E., & Brancik, K. (2021). Services for zero trust architectures—a research roadmap. *2021 IEEE*

- International Conference on Web Services (ICWS)*, 14–20.
- Buck, C., Olenberger, C., Schweizer, A., Völter, F., & Eymann, T. (2021). Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust. *Computers & Security, 110*, 102436.
- Campbell, M. (2020). Beyond zero trust: Trust is a vulnerability. *Computer, 53*(10), 110–113.
- Chen, Y., Hu, H., & Cheng, G. (2019). Design and implementation of a novel enterprise network defense system by maneuvering multi-dimensional network properties. *Frontiers of Information Technology & Electronic Engineering, 20*(2), 238–252.
- Compastié, M., Badonnel, R., Festor, O., He, R., & Kassi-Lahlou, M. (2016). A software-defined security strategy for supporting autonomic security enforcement in distributed cloud. *2016 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, 464–467.
- Cunningham, C. (2018). The zero trust extended (ztx) ecosystem. *Forrester, Cambridge, MA*.
- Cunningham, C. (2020). A look back at zero trust: Never trust, always verify. *Forrester*.
- Garbis, J., & Chapman, J. W. (2021). *Zero Trust Security: An Enterprise Guide* (1st ed.). Apress.
- Gardner, W., Mulvey, E. P., & Shaw, E. C. (1995). Regression analyses of counts and rates: Poisson, overdispersed Poisson, and negative binomial models. *Psychological Bulletin, 118*(3), 392.
- Kerman, A., Borchert, O., Rose, S., Tan, A., & others. (2020). Implementing a zero trust architecture. *National Institute of Standards and Technology (NIST)*, 75.
- Kindervag, J., Balaouras, S., Mak, K., & Blackborow, J. (2010). *No more chewy centers: The zero trust model of information security*. Forrester Research.
- Mehraj, S., & Bandy, M. T. (2020). Establishing a zero trust strategy in cloud computing environment. *2020 International Conference on Computer Communication and Informatics (ICCCI)*, 1–6.
- MITRE Corporation. (2024a). *MITRE ATT&CK* [Dataset]. <https://attack.mitre.org>
- MITRE Corporation. (2024b). *MITRE D3FEND* [Dataset]. <https://d3fend.mitre.org>
- Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture* (NIST Special Publication 800-207). National Institute of Standards and Technology (NIST). <https://csrc.nist.gov/publications/detail/sp/800-207/final>
- RStudio Team. (2020). *RStudio: Integrated Development Environment for R*. RStudio, PBC. <http://www.rstudio.com/>
- Samani, R., Honan, B., & Reavis, J. (2014). *CSA guide to cloud computing: Implementing cloud privacy and security*. Syngress.
- Tyagi, M., & Arora, K. (2022). *Benefits of Zero Trust Adoption Through the Lens of MITRE ATT&CK and D3FEND Frameworks*. F5. <https://www.f5.com/resources/reports/office-of-the-cto-zero-trust-and-the-mitre-frameworks>
- Verizon. (2024). *The VERIS Community Database (VCDB)* [Dataset]. <https://github.com/vz-risk/VCDB>
- Verizon Enterprise Solutions. (2024). *2024 Data Breach Investigations Report*.
- Ward, R., & Beyer, B. (2014). Beyondcorp: A new approach to enterprise security. ; ; *Login: The Magazine of USENIX & SAGE, 39*(6), 6–11.
- Yan, X., & Wang, H. (2020). Survey on zero-trust network security. *Artificial Intelligence and Security: 6th International Conference, ICAIS 2020, Hohhot, China, July 17–20, 2020, Proceedings, Part I 6*, 50–60.
- Yeoh, W., Liu, M., Shore, M., & Jiang, F. (2023). Zero trust cybersecurity: Critical success factors and a maturity assessment framework. *Computers & Security, 133*, 103412.
- Zaheer, Z., Chang, H., Mukherjee, S., & Van der Merwe, J. (2019). eztrust: Network-independent zero-trust perimeterization for microservices. *Proceedings of the 2019 ACM Symposium on SDN Research*, 49–61.

8. Acknowledgment

This work was supported in part by U.S. Army Combat Capabilities Development Command (DEVCOM); U.S. Army Research Laboratory (ARL); ARL South at the University of Texas at San Antonio (UTSA) and the Army Educational Outreach Program (AEOP).