Comparative Analysis of Incident Reporting Frameworks Against MITRE ATT&CK

Matthew Ryan[a]

[a]*Beacom College of Computer and Cyber Sciences, Dakota State University, Madison, 57042, SD, USA*

Abstract

This study concentrates on the accuracy and efficacy of three selected incident reporting schema as assessed against various techniques defined within the MITRE ATT&CK Framework, identifies specific strengths and gaps within each framework, and extracts findings from this assessment to propose a universally usable incident reporting schema or taxonomy. The study intends to identify if there are opportunities to improve the initial reporting of a cyber incident from the perspective of accuracy and clarity. Various organizations approach initial reporting from differing angles, and as of the date of this study, there is arguably no industry standard for cyber incident reporting. As a result, it is possible that organizations or governmental entities may be utilizing incident reporting schema that can result in confusion, conflation, or inaccurate reporting. This study sought to identify these potential inconsistencies and gaps in three large-scale incident reporting schema by notionally mapping detectable adversarial hacking techniques as reportable events within each schema. Using clustering, consistent patterns were identified for analysis and derived insights, specifically potential confusion regarding how certain MITRE ATT&CK techniques would be categorized in each of the taxonomies. For instance, whether a technique should be categorized as "Malicious Logic" or "User Level Intrusion," "Web Application Attack" or "System Intrusion," and "Web vector" or "Email/Phishing." These findings were utilized to develop a new incident reporting taxonomy that utilizes the strengths of each assessed framework as a universal and publicly available taxonomy. The Universal Cyber Incident Taxonomy (UCIT) features a simplified and hierarchical approach to categorizing cyber events similar to Linnaean Taxonomy for biological organisms. By defining individual events with a leveled categorization system, events are more accurately described,

July 10, 2025

and there is less opportunity to conflate similar events or mislabel them.

Keywords: Cyber Incident Reporting, MITRE ATT&CK, CJCSM6510.01B, Verizon DBIR, CISA Incident Reporting, Cyber Incident Response, cybersecurity

---

1. Introduction

As cybersecurity defensive operations mature across organizations, the ability to document, categorize, and report cyber incidents has become another facet of the cyber lexicon, along with the strengths and challenges that come with it. The term cyber attack generally refers to criminal activities conducted via the Internet, typically to steal money and confidential information or to advance a political or ideological goal (hacktivism) [1]. As attacks persist within organizations, a need exists to properly and adequately categorize and document these incidents for the record, either for archival or upstream reporting. However, due to individual mission needs and requirements, incident reporting schema has been primarily organization-based and customized in lieu of a single standardized framework provided by an organization like the National Institute of Standards and Technology (NIST). Indeed, NIST has published cyber incident reporting and response guidance [2], but this report merely outlines critical data fields that should be included in any cyber incident report. Further, NIST SP 800-61 does not provide any guidance or standardization for the defining naming conventions an organization could use to refer to its cyber incident reports when categorizing them for further reporting, retrospective analysis, or archival. The frequency of attacks faced by the average host connected to the Internet remains elevated [3]. Therefore, organizations must define and deploy their own incident categorization schema, which tends to be decentralized, individually focused, and entirely non-standard.

As a result, it appears as though nearly all organizations have defined their own unique incident reporting schema with high-level incident "categories" that assist in metrics and statistical collection for retrospective meta-analysis and reporting [4]. Within a Security Operations Center (SOC) environment, these categories are generally how a detected event comes to be defined throughout the incident handling process and can either change as new information is discovered, potentially result in duplicate reporting as new event data is detected, or kept permanently even if the understanding

of the event changes. The diversity of schema yields an opportunity for comparative analysis. Several of these incident reporting frameworks are likely imperfect systems that can confuse an analyst regarding proper event identification, events potentially overlapping multiple categorization schema, or events not truly fitting in any of the named categories. If this is the case, this results in inaccurate data sets for large-scale organizations. When combined with a lack of uniformity across organizations, it results in bluedata sets that do not correlate.

When there is a lack of aforementioned mandated uniformity, part of the difficulty is a similar lack of requirement to make these schema publicly known [5]. This severely limits the sample population of incident reporting schema that can be used for such a study. Initial research suggests that incident reporting frameworks that include actual category definitions are relatively sparse. Most organizations likely treat such schema as internal documentation, with no inherent responsibility to make them publicly available. Fortunately, three extremely vast organizations make their schema available and include category definitions. The Department of Defense (DoD) mandates that all cyber incidents and events be reported in accordance with the guidance published by the Chairman of the Joint Chiefs of Staff Manual (CJCSM) 6510.01B Cyber Incident Handling Program [6]. Similarly, within the United States government, the Cybersecurity and Infrastructure Security Agency (CISA) mandates any incident reporting to US-CERT follow the CISA Incident Reporting Guidelines [7]. Finally, multinational telecommunications conglomerate Verizon releases an annual Data Breach Investigations Report (DBIR) [8] that includes incident/event category information and definitions. These three frameworks were subsequently selected simply due to their availability combined with their large-scale size and breadth across their stakeholders and/or customer base.

This research study aims to compare and contrast these three publicly available incident reporting schema by selecting known adversarial hacking events from the industry standard MITRE ATT&CK framework [9], following the guidance of each incident reporting schema and categorizing each event appropriately. The research focuses on identifying ambiguity within publicly available incident reporting taxonomies and utilizing findings to propose a new universal taxonomy. Conducting the study, example techniques were selected from each of MITRE ATT&CK's 14 technique families. Example events were categorized according to the guidelines within each organization's reporting schema and analyzed for potential gaps, du-

plication, or insufficient definitions. Following an exploratory assessment of these techniques, a midpoint assessment report was prepared, and findings subsequently determined the evaluation method was relatively sound. The selection of techniques from the ATT&CK families was subsequently doubled for the sake of strengthening the assessment data set, and incident reporting taxonomies were reassessed. This also provides opportunities for future work, as the total number of MITRE ATT&CK techniques at the time of this writing currently stands at 466 techniques, and 86 techniques were assessed.

Desired insights from the analysis include identifying gaps in certain schema, areas where poor definitions result in gaps or duplication, and the ability of the organization to retrospectively review past results for accurate statistical analysis for the purposes of proficiency monitoring or resource allocation. Findings surrounding logic errors in accurate event labeling were identified, and notable strengths within each evaluated taxonomy were identified. These findings were then utilized to create a new taxonomy derived from each taxonomy's identified strengths and weaknesses. The end state of this research is the development of the Universal Cyber Incident Taxonomy (UCIT), designed to be publicly available, freely usable, and universally operable for any public or private sector SOC. It is currently being tested in a federal government SOC in parallel against an internal taxonomy for further comparative analysis.

As an end-product of this analysis, UCIT is designed as a modular, publicly available, and freely usable cyber incident taxonomy that addresses all potential weaknesses from the assessed incident reporting schema. It derives the inherent strengths from each taxonomy. It utilizes taxonomy based on what is used for biological organisms (Kingdom, Phylum, etc.) to accurately and adequately categorize incidents based on a ranked hierarchy of data fields. UCIT was then subjected to the same cross-walk evaluation as the other taxonomies for comparative analysis purposes.

The contributions of this work thus include:

- An assessment and evaluation of three large incident taxonomies currently in operational use today.
- An analysis of potential causes of confusion, conflation, or miscategorization due to inefficient or incomplete incident category definitions.
- A proposed incident taxonomy specifically created based on these lessons learned and freely released for adoption and use.

The remainder of the paper is structured as follows. Section 2 discusses related work and other efforts associated with the proper incident definition and categorization. Section 3 describes the mechanism for testing a number of selected incident taxonomies and the data set used to test them. Section 4 discusses the outcomes of the experiment's notional incident mapping. Section 5 is a dissection of some of the identified challenges that appeared consistent across all assessed taxonomies blueand a few identified strengths. Section 6 leverages the prior analysis as its basis for conceptualizing and developing a new framework based on lessons learned. Finally, Section **??** identifies areas where more research work could be conducted to advance the concept, as well as how UCIT as a framework could be further tested.

2. Related Work

Many studies specifically concentrate on cyber incident reporting, and obstacles surrounding accurate and meaningful reporting are well-researched and understood. Authors in [10] comparatively analyzed five different reporting templates from a variety of European cybersecurity organizations, though they concentrated specifically on necessary data elements for robust reporting. It did not, however, concentrate specifically on the challenges that come with incident categorizations as a mechanism of organizational sorting of incidents/events. Insights from this study resulted in recommendations for additional standardized fields and a proposed template write-large rather than a specific categorization schema.

The work reported in [11] takes a similar approach, selecting alternative incident reporting schema such as the industry-standard STIX and highlighting the strengths and weaknesses of the reporting formats of six different schemas. However, the interesting difference between this work and [10] is how incident reporting formats were defined. While [10] concentrated on organizational reporting for higher level awareness and command/control, [11] defined incident reporting as "information sharing" and selected schema specifically designed to highlight the sharing of cyber event technical indicators more in line with concepts of the cyber threat intelligence field. Specifically, Structured Threat Information eXpression (STIX) and the other evaluated frameworks within this study are explicitly geared towards highlighting technical indicators for information sharing with a large community so that the community can hunt for similar activity within their environments. While elements of this information likely exist in all cyber incident report-

ing schema, the overarching intent behind STIX and similar frameworks is to provide a standardized data schema for the machine-ingest of network and host-level data rather than a summary narrative-based descriptions or categorizations of the event itself.

Lif et al.'s work [12] is most closely aligned with the intent of this research proposal, in which 16 different cyber defense organizations were asked to file cyber incident reports using their schema. At that time, they were compared to pre-defined critical information elements and assessed on accuracy and maturity in reporting. This work is closely aligned with this proposal thematically but does not concentrate specifically on top-level incident categorizations. Organizations were assessed based on how well they reported information elements within their schema, not how accurately their categorization method was for the overarching event. As an additional minor note, all researched publications also concentrated solely on European organizations, except for Menges et al.'[11] concentration on STIX, as aforementioned. It appears as if no American cybersecurity organizations were utilized in a comparative analysis. While this is arguably irrelevant to the wider research space in this area, it is likely beneficial to assess American governmental and industry organizations for the benefit of improving American organizational incident reporting. Additionally, comparative analysis between American and European organizations (or any other international organizations) is likely an area where future research may be beneficial.

Following the completion of the study attempts to identify related work were revisited as a means to find additional taxonomies for comparative analysis potentially. Zaccaro et al. [13] appears deceptively similar, as this study concentrated on a multilevel approach to a cyber security taxonomy. However, this study is specific to incident response performance, either by an individual analyst or a SOC. The study specifically develops a multilevel taxonomy for measuring performance and efficiency during incident handling. This differs from a taxonomy for specific categorization of cyber incidents.

Additionally, a large number of preexisting research concentrates specifically on intrusion detection rather than the necessary subsequent intrusion reporting. Yuill et al. [14] propose a military battlefield-intelligence process approach for intrusion detection in cybersecurity environments. This is a similar research area and gap but does not address the specific challenges of effective reporting and categorization of incidents post-compromise or in progress.

Attempting to identify if highly specialized subsets of cybersecurity could

be expanded to the field writ-large, Zhu et al. [15] was reviewed to see if specific categorization insights could be utilized within this study. Unfortunately, this study specifically concentrates on identifying a taxonomy for cyber attacks on SCADA systems. This particular sub-field is too explicit in scope to apply to large-scale cybersecurity organizations seeking to accurately and adequately classify cyber incidents via a taxonomy or schema. Indeed, further research may be warranted to identify if specific and highly specialized fields within cybersecurity deserve their taxonomy due to their uniqueness.

Finally, Ibrishimova's work [16] is similar to this study, as it recognizes that the cyber threat landscape is changing rapidly, thus making the process of scientific classification of incidents for incident response management difficult. This is indeed factual and related to the intended concept of this study. However, the researchers specifically concentrate on how existing efforts to automate the process of incident classification do not make a distinction between ordinary events and threatening incidents, which can cause issues that permeate throughout the entire incident response process. This differs from the scope of this study, as the researchers propose a machine-learning model to detect the probability of malice in a given event.

3. Methodology

In simple terms, given the dual assumption that all SOCs operate some sort of labeling scheme for their identified cyber incidents and events and that said schemes are organizationally specific and disparate, this study intends to identify a selection of large-scale categorization taxonomies, assess them against various cyber-attacks, and specifically identify any weaknesses in said taxonomies, where there may be notable strengths of one against another, and gain insight into how an incident categorization schema can be improved or matured. Considering prior work has primarily concentrated on the assessment of taxonomy and broader data elements, opportunities exist to understand standard categorization of incidents and events in large cyber defense organizations, best practices, and potential challenges that come with attempting to utilize a categorization system for labeling and sorting cyber incidents & events at scale. Conducting such a study essentially requires three data collection stages.

First, a large sample size of diverse categorization schema is preferred, but as aforementioned, very little data is publicly available. In lieu of mul-

tiple schema, arguably, the next best sample set is the best effort for large organizations that cover large swathes of network space. Secondly, a relevant and reliable catalog of cyber-attacks and offensive actions is required to assess each of the selected schemas. This catalog must be robust enough to capture all known potentially detectable activity that could subsequently be used for filing a cyber incident report. Only from there is a viable cross-walk possible, in which certain cyber activities can be sorted and organized based on each reporting schema incident category definitions. In other words, this research will identify example cyber events for usage in a notional categorization exercise within the constraints of each selected incident reporting schema. From there, the analysis would involve a comparison of various categorization schemas to identify best practices, gaps, and weaknesses across all selected frameworks.

3.1. Selecting the Frameworks

As mentioned, research work specifically focused on collecting multiple incident categorization schema was largely unsuccessful, as very few organizations publicize their preferred taxonomy and labeling, much less communicate whether any standardization with other known schema is occurring. Subsequently, "best effort" sample collection occurred primarily concentrating on large organizations with well-known and published schema, notably the aforementioned DoD's CJCSM 6510.01B, Verizon's Incident Classification Patterns, and CISA's US-CERT Federal Incident Notification Guidelines. These schema are some of the few that are publicly available and have internet access but also benefit from their hosting organization's size. Exercising these schemas should result in a relatively sizable impact on the research findings based on the number of systems and frequency of reporting each of these organizations likely participate in. A challenge within this research area is finding a categorization schema that is publicly available and not customized to the individual organization. Still, these three organizations are likely large and mature enough to ascertain relevant insights from their event/incident categories.

Within the Department of Defense, CJCSM 6510.01B is the governance document for incident handling, major processes, and hierarchical reporting of cyber incidents within the entirety of the DoD Information Network (DODIN). It documents the workflow for incident reporting and response and inter-departmental relationships for information sharing and reporting. Most importantly and relevantly, however, is that CJCSM 6510.01B defines the 9

incident categories all reported cyber events must fall into to be reported to USCYBERCOM. These 9 categories are numerical, each corresponding to a specific incident behavior. Notably, Root level compromises, User level compromises, Unsuccessful Attempts, Denial of Service, Noncompliance activity, Reconnaissance, Malicious Logic, Explained Activity, and "Investigating" to denote newly opened events. CJCSM 6510.01B also includes instructions for incident handling, particularly regarding the shifting of incident categories as more information is understood [6]. An example scenario, for instance, may involve an analyst opening a report for "Investigating," later re-categorizing it as "Malicious Logic," before eventually labeling it as "Root Level Compromise" to complete the incident handling workflow. The latest iteration of CJCSM 6510.01B was released in 2012 and is still in use by Defensive Cyber Operations (DCO) analysts and operators to this day across the DoD. An initial review of these incident categorizations suggests opportunities for analyst confusion and multiple interpretations of definitions. For instance, "Malicious Logic" as an incident category may or may not also involve "User Level Intrusion," and a number of different hacking attacks may appropriately align to either or both categories. Should this opportunity for multiple categorizations bear out within the assessment, it should make for a particularly valuable finding for this particular framework, if not large. The various CJCSM6510.01B [6] event/incident categories are summarized in Table A.8 (see Appendix A).

Verizon's Security Operations Center is extremely large, considering its position as an Internet Service Provider (ISP). It provides a unique look into an extremely large data set that is publicly available and published annually in its *Data Breach Investigations Report (DBIR)* [8]. This public disclosure also allows public insight into how Verizon categorizes and sorts its incident data in a way that can be similarly assessed in this proposed study. These categories primarily revolve around high-level explanations of the threat landscape. It's also important to note that these categories may change yearly. Verizon appears to utilize high-level definitions for the attack style. It relies on 8 category titles, including Social Engineer, Basic Web App Attacks, System Intrusions, Miscellaneous Errors, Privilege Misuse, Lost/Stolen Assets, Denial of Service, and "Everything Else" [8]. Notably, there is no publicly available policy for identifying how an activity may be re-labeled, which may be a challenge within the research. Also, there is no "Investigating" style status to denote that an event is being analyzed; notionally, all data is completely analyzed at the time of publication. It is unclear if the schema

published within their annual DBIR is used within Verizon's internal security operation and DCO construct. However, this schema was selected primarily for its simplicity, clarity, and many incidents categorized using it. The Verizon DBIR [8] incident reporting schema is summarized in Table A.9 (see Appendix A).

Finally, CISA's Federal Incident Notification Guidelines [7] provide the categorization schema for all government agencies (outside of DoD) to report through in accordance with CISA's federal oversight mandate. This particular schema is interesting because CISA has the comparatively unique challenge (compared to other selected schema) in identifying a categorization scheme that all reporting private industries and government agencies can adhere to despite their own individually unique missions, networks, and likely customized incident reporting formats. As a result, it appears as though CISA has attempted to require only the most critical of elements in its reporting schema. However, it still involves fields that are used for sorting and categorization. Notably, the primary table within CISA's schema is "Attack Vectors Taxonomy," in which various examples and descriptions are provided for various named attack vectors. CISA utilizes 9 categories, including Unknown, Attrition, Web, Email/Phishing, External/Removable Media, Impersonation/ Spoofing, Improper Usage, Loss or Theft of Equipment, and "Other." Indeed, CISA appears to employ several required data fields. Still, it appears as though "Attack Vectors" appears to exist as the one consistent mechanism of any sort of simple categorization schema. The Attack Vectors categories and definitions [7] are summarized in Table A.10 (see Appendix A).

Each of the selected schemas represents a very large organization with a robust DCO detection, reporting, and response mechanism, each by the necessity of its purpose and mission. While follow-on work to this study will always benefit from the identification and inclusion of more schema, the currently publicly available labeling frameworks appear to be limited to these three large organizations.

3.2. The MITRE ATT&CK Framework

Assessing each schema for accuracy and clarity is only as effective as the data inputs fed into each framework. By extension, the diversity of attack vectors similarly exercises each schema to its fullest capability, maximizing opportunities for logic error, interpretation conflicts, and other potential breakdowns in the provided definitions from each organization. As a result,

to conduct a comparative analysis of the categorization schema, there is a need for national cyber incident data to cross-walk the frameworks. MITRE ATT&CK, or *The Adversarial Tactics, Techniques, and Common Knowledge (MITRE ATT&CK)*, is a guideline for classifying and describing cyberattacks and intrusions. It was created by the MITRE Corporation [9] and released in 2013. As the accepted industry standard for categorizing various types of cyber events and activities, there are arguably few better options short of a cyber range exercise for assessing the accuracy and maturity of the schema as mentioned above from an incident categorization standpoint.

MITRE ATT&CK divides various malicious cyber activities into 14 technique families, each with as few as 7 or 42 individual techniques known to be utilized in cyber-attacks [9]. These technique families are defined in Table A.11 (see Appendix A). These techniques are typically chained together into playbooks and can be used to typify and describe various threat actors, nation-state actors, or other hacking groups [17]. Notionally, any one technique is detectable, and MITRE similarly documents detection techniques for each event, if applicable, in a way that a SOC analyst would be able to reference and mark certain techniques appropriately by name or technique labeling convention. Essentially, MITRE ATT&CK serves as an attack encyclopedia in which all known techniques used to conduct offensive cyber activity are sorted, categorized, and documented. [18]

Additionally, many techniques are further contextualized into various subtechniques for greater granularity and detail. This study randomly selected three different attacker techniques from each of the 14 MITRE ATT&CK attack families for diversity in testing. These families include Reconnaissance, Resource Development, Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, Collection, Command and Control, Exfiltration, and Impact.

3.3. Application of Assessment Cross-Walk

The assessment methodology for the categorization schema revolves around utilizing each in a notional incident reporting scenario, in which the researcher maps the MITRE ATT&CK technique, based on how MITRE defines it, against the most appropriate category or categories within each organization's labeling convention. Fundamentally, the assessment mimics the decision tree of a SOC analyst functioning in each of these organizations, in which a detection of the specific ATT&CK technique occurs and is sorted appropriately based on the definitions of each organization's categorization

schema. In other words, the selected techniques will be "filed" as incident reports and tagged with appropriate incident categories for each schema. Special consideration will be made for techniques that may not apply to the definitions of any one incident category, as well as those that may apply to multiple incident categories, for the purposes of identifying gaps. These techniques will be plotted on each organization's categorization schema matrix (see Section 4) by cross-referencing the ATT&CK definition with the most closely aligning category or categories for each organization.

The end-state of this phase is to have multiple categorization schema matrices completed with various MITRE ATT&CK techniques recorded and plotted appropriately against predefined incident categories for each of the three organizational taxonomies. To provide a notional example, a detection of the MITRE ATT&CK technique of "Access Token Manipulation" would result in an incident category of "CAT1 - Root level Compromise" for DoD, "System Intrusion" for Verizon, and "Unknown" for CISA within each of their incident categorization schema. Since this is a technique within the "Privilege Escalation" ATT&CK family, it would result in compromise at the network level, which aligns with the appropriate definitions for Verizon and DoD. Still, the initial attack vector is not identified for US-CERT's reliance on the "Attack Vector," which results in these categorizations respectively. Broadening the scope, this would then be applied to selected ATT&CK families and techniques across the entire ATT&CK framework for the purposes of wider and more robust data collection of categorization accuracy and efficacy.

Sub-technique selection was a matter of balancing data diversity with available time. Upon cursory review of the MITRE ATT&CK sub-techniques in totality, it was identified that different values had significantly different "Procedure Examples" attached to each data record. Procedure examples are documented real-world instances of the sub-technique as observed in technical writeups, articles, or breakdowns of offensive activity [19]. The initial selection consisted of the top 3 sub-techniques with the most procedural examples per each parent technique. This was intended to represent the most commonly observed Sub-techniques in a typical operating environment based on the amount of available documentation on each sub-technique [17]. The total was later doubled to the top 6 sub-techniques to adequately increase diversity in the assessment data set.

Overall, 84 total techniques spread across all 14 ATT&CK families (6 per family) were selected. During the midpoint assessment period, 42 techniques were identified with evidence of insightful clustering of incident labels sur-

rounding certain techniques, both vertically and horizontally. Vertical clustering suggests multiple categories within a taxonomy apply to the given observed event, which suggests duplication and redundancy, which is a negative trait considering how this would result in confusion from the filing analyst. Horizontal clustering was considered a positive trait, indicating consistent definitions applying to multiple observed events within an attack path. In other words, the category is well defined if multiple actions within a "System Intrusion" could accurately result in a consistent "System Intrusion" label.

3.4. Analysis Phase

Following the Assessment Cross-Walk, pattern analysis, primarily through clustering, was employed to derive insights into any strengths and weaknesses of each respective categorization taxonomy. It was postulated that consistent gaps in category definitions would be identified due to this clustering, particularly if partially effective or incomplete definitions existed for the techniques "detected." For instance, by limiting initial categorization to "Attack Vector," such as in the case of US-CERT, actions detected further down the cyber kill chain would need to be sorted as "Unknown" as it was not immediately apparent how initial access into the network was accomplished by the notional attacker. Similarly, another expected result was that simplicity would be the preferred tactic for initial categorization, such as in the case of Verizon's simple "System Intrusion" category, which does not discriminate based on granular details of the detected technique and labels the action as an "intrusion" without much more description. This also suggests overly descriptive detail can confuse the part of the filing analyst, where multiple categories of the incident may apply to an individual technique, such as in the case of the DoD CJCSM 6510.01B guidance. Specifically, from a technical standpoint, a detected technique may be accurately categorized as "Malicious Logic" and "User Level Intrusion," which results in confusion in definition interpretation and would notionally lead to inaccurate data sets depending on the opinion of the individual analyst.

Finally, the findings portion of the study will concentrate on categorization best practices and recommendations for accurate reporting in any categorization schema. Ideally, findings would be robust enough to provide a new and improved notional categorization schema for any organization moving forward. The intent of the study ultimately is to identify the best ways to provide a quick and summative term or definition of a cyber event that accurately typifies the observed technique for the purposes of accurate reporting,

retrospective data analysis of incident statistics, and guidelines for curating a comprehensive data-set of recorded cyber incidents that can be utilized for other business purposes. This would include many use cases, investment in additional detection technologies, human resource allocation, etc. Out of each of these three large organizations' categorization schema, best practices, and recommendations should be able to be identified that would notionally feed recommendations for improvements or a new consistent and uniform categorization taxonomy outright.

3.5. Limitations to this Approach

The limitation to this assessment approach is that it is extremely subjective, prone to researcher bias and potential error. However, the suspected ambiguity in how the subject organizations categorize their cyber incidents is similarly just as subjective, relying on the judgment of individual analysts working within their SOCs. Assuming the provided incident category definitions are the measure for which an analyst judges an observed event and no other unpublished guidance within each organization exists, the categorization ambiguity is quickly apparent without any in-depth assessment. The cursory review would suggest that any number of techniques, when observed by a filing SOC analyst, would result in confusion and ambiguity as the analyst forces unique events into imperfect category descriptors. Thus, this assessment attempts to simulate a "day in the life" of a typical SOC analyst working within all of these organizations, categorizing incidents as they are observed in logs and artifacts. Future work utilizing large sets of test users to cluster categorization trends further would likely strengthen this argument since it is highly suspected that categorization ambiguity will still occur.

The assessment was conducted by one of the authors, a SOC analyst with 15 years of experience working entirely within 24/7 Federal and DoD SOC environments, in which he observed, reported, and responded to incidents within 2 of the 3 assessed frameworks. This assessment was then peer-reviewed and approved by a senior colleague, a SOC analyst with 23 years of experience entirely within the same environments. Categorizations were determined by analyzing the sub-technique's definition and applying all relevant and possible organizational categorizations to them, even if overlap existed, per the definition provided by each organization. For example, if little definitional difference existed between what constitutes a DoD rating for "CAT2" vs. "CAT7" via the assessed sub-technique, both were selected. Similarly, if the schema had little relational relevance to the sub-technique in

question, such as much of CISA/US-CERTs schema regarding Attack Vector, all possible iterations of how that sub-technique could have been observed were selected or were labeled "Unknown" if all possible attack vectors could logically lead to the assessed sub-technique eventually occurring.

As aforementioned and later, these frameworks and the resulting proposed taxonomy are currently being assessed within an operational Federal SOC. This is intended to address the aforementioned limitation regarding researcher bias and test user population. This is currently ongoing and is suspected to be met with substantial time delay as it is not only completed but also must undergo an information security review before it is approved for any public dissemination. Future work is planned to document this assessment and follow up on this research.

4. Findings and Results

In this section, some important key findings and results are described. In the first phase of this study, 42 techniques were selected, three from each of the 14 ATT&CK families, for the purposes of diversity in testing that spans across all steps of the cyber kill chain. During the assessment mapping, the detected technique was identified and assumed "detected" by a notional cyber analyst and filed according to each of the selected organization's most appropriate cyber incident category labels or labels. To mimic a true operational environment most closely, each technique was assumed to be detected accurately and individually. While it is likely that any individually detected technique would logically result in a follow-on retrospective analysis of surrounding events by a competent analyst, for the purposes of this study, the assumption is that the detection would be filed immediately (and subsequently labeled immediately) to hit key performance indicators (KPIs) and required reporting time constraints. This generally aligns with the required timelines within most incident reporting frameworks, specifically the general guidance that a report should be filed as quickly as possible with the intention that additional information will be added as it is identified.

4.1. Reconnaissance

Moving through each ATT&CK family [9], the first phase is Reconnaissance, which consists of techniques that involve adversaries actively or passively gathering information that can be used to support targeting. This can include vulnerability scanning, searching the internet for access points or

Table 1: Reconnaissance, Initial Access, and Execution.

		Reconnaissance						Initial Access						Execution					
		Active Scanning	Search Open Websites/Domains	Gather Victim Identity Information	Gather Victim Org Information	Search Victim-Owned Websites	Search Open Technical Databases	Drive-by Compromise	Exploit Public-Facing Application	Phishing	Supply Chain Compromise	Trusted Relationship	Hardware Additions	User Execution	Command and Scripting Interpreter	Exploitation for Client Execution	Windows Management Instrumentation	Native API	System Services
CISA / US-CERT	Unknown	X	X	X	X	X	X				X			X	X	X	X	X	X
	Attrition																		
	Web	X	X	X			X	X	X		X	X			X	X	X	X	X
	Email/Phishing									X		X		X	X	X			
	External/Removable Media									X			X						
	Impersonation/Spoofing							X						X	X	X	X		X
	Improper Usage								X										
	Loss or Theft of Equipment												X						
	Other	X	X	X							X	X							
DOD CJCSM 6510.01B	CAT-1 Root Level Intrusion							X	X	X	X	X	X	X	X	X	X	X	X
	CAT-2 User Level Intrusion							X	X	X	X	X	X	X	X	X	X	X	X
	CAT-3 Unsuccessful Activity																		
	CAT-4 Denial of Service																		
	CAT-5 Non-Compliance Activity										X		X						
	CAT-6 Reconnaissance	X	X	X	X	X	X												
	CAT-7 Malicious Logic							X	X	X	X	X	X	X	X	X	X	X	X
	CAT-8 Investigating																		
	CAT-9 Explained Activity																		
Verizon DBIR	Basic Web Application Attacks								X					X		X			
	Denial of Service																		
	Lost and Stolen Assets												X						
	Miscellaneous Errors																		
	Privilege Misuse											X							
	Social Engineering									X				X					
	System Intrusion							X	X	X	X	X	X	X	X	X	X	X	X
	Everything Else	X	X	X	X	X	X												

leaked credentials, or profiling potential targets and their infrastructure. As shown in Table 1, immediately apparent is how DoD specifically quantifies this activity with its incident category (CAT-6 Reconnaissance). In contrast, CISA/US-CERT relies on labeling by Attack Vector, and Verizon does not seem to consider this activity an incident at all.

4.2. Resource Development

Along those same lines, the ATT&CK family of "Resource Development," or techniques that involve adversaries creating, purchasing, or compromis-

ing/stealing resources that can be used to support targeting. Since this involves weaponization and staging of attack infrastructure, it appears as if none of the organizations specifically identify this activity as "incident worthy," and none of the selected techniques explicitly map to any category definitions for any organization. This is likely because this stage involves no active interaction with the organization's network and is not considered an "active attack." MITRE even categorizes each of these technique families as "PRE-attack" within multiple fields of its data set, including the "Mitigation" field and "Platform" field.

4.3. Initial Access

As anticipated, "Initial Access" yielded the most variation in categorization clustering, particularly within the CISA/US-CERT construct, primarily due to its concentration on "Attack Vector" as its primary categorization mechanism. As shown in Table 1, Initial Access consists of techniques that use various entry vectors to gain their initial foothold within a network. Techniques such as Drive-By Compromise, Exploitation of Public-Facing Web Servers, and Phishing all involve various Attack Vectors depending on how the technique is deployed against a target network. While this works particularly well for the CISA/US-CERT construct, first signs of interpretive confusion are apparent within the DoD labeling scheme, as any individual technique could be labeled "CAT-7 Malicious Logic" ("Installation of software designed and/or deployed by adversaries with malicious intentions to gain access to resources or information without the consent or knowledge of the user [6]" but also "User Level" or "Root Level" compromise simultaneously. Similar clustering exists within the Verizon construct as well, as this activity could be identified as a "System Intrusion" simultaneously with a "Basic Web Application Attack" or "Social Engineering," depending on the nature of the technique. However, this does not appear to be a case of overlapping definitions so much as a mechanism to differentiate between different Initial Access techniques. For instance, phishing delivering malware successfully would be identified as a "System Intrusion," phishing impersonation fooling a user into divulging credentials would be notionally labeled "Social Engineering."

4.4. Execution

The ATT&CK family "Execution" (see Table 1) consists of techniques that result in adversary-controlled code running on a local or remote system.

Techniques that run malicious code are often paired with techniques from all other tactics to achieve broader goals, like exploring a network or stealing data. Each organization's labeling schema appears to have the same clustering dependent on the specific technique deployed. Within this ATT&CK family, the attack vector may differ, as Execution is essentially the logical simultaneous step as Initial Access. In other words, successful malware delivery via phishing would immediately result in User Execution of malicious code in the case of CISA/US-CERT. Definition overlap appears to persist within the DoD construct, as the CAT-1/2/7 definitions appear to apply to all of the listed techniques, resulting in further data inaccuracy and analyst confusion. Notably, the Verizon construct appears to begin its trend that is continued throughout the other ATT&CK families, where categories begin to divulge into a single "System Intrusion" label consistently. This may be a positive or negative trait for usability, pending additional analysis.

4.5. Consistency Amongst the Remaining Families

From here, the remaining ATT&CK families appear to result in consolidation in categorization across all three organizations' categorization schema without significant variation (see Tables 2, 3, and 4, respectively). In this portion of the cyber kill chain, the attacker is already on the network, and detections would similarly align to catching an adversary executing activities within the target infrastructure. As a result, the consolidation of categories and subsequent horizontal clustering makes sense, as there is less diversity in interpretation regarding what exactly is happening within an individually detected technique. Notably, detecting these techniques so far down the cyber kill chain seems to bypass the CISA/US-CERT categorization scheme entirely, as the attack vector at this point of an intrusion would be entirely unknown. An analyst would then have to work backward to identify where initial access occurred and subsequently identify an initial attack vector. DoD categorization remains consistently overlapped, as the CAT-1/2/7 issue persists, wherein an analyst can technically provide an accurate categorization using any of these labels mentioned above. As noted, individually, this is likely not problematic. Still, given a large data set of incidents, it would potentially result in the same techniques being categorized differently depending on the event and the filing analyst, which is an issue for maintaining an accurate data set. Verizon's construct defaults to "System Intrusion" for nearly the entirety of the remaining listed techniques, which is indeed accurate. However, further analysis is needed to determine if this simplis-

tic approach is preferred or if additional categories that typify the activity result in a more accurate categorization scheme. This cluster patterning across each organization is also notable, as it appears to provide consistent findings for the remainder of the techniques within the cyber kill chain that will likely provide best practices or recommendations for improvement when the analysis is completed.

4.6. Expanding the Sample Size

Within the second evaluation phase, 3 additional techniques were added to every family to increase the detail of the data set and effectively double the size of sample events for assessment. This increased the number of sampling events to 86, with 6 per each of the 14 technique families. This was primarily done to provide greater variation to the data set to rule out any potential for false positive trending via coincidental correlation of event clusters. This did not prove easy to select variations of techniques, considering that the techniques were sorted into higher summations of technique families, so these events are highly similar. That said, effectively doubling the data set's size further strengthened the identified findings. Additional techniques aligned almost entirely with expectations developed via earlier evaluations. No significant deviations from preexisting clusters surfaced, suggesting that the current findings sustained further evaluation. Generally speaking, vertical clustering remained consistent within the wider Technique Families, such as Initial Access. Horizontal clustering across multiple Technique Families again supported the assertion that the categorization label was sufficiently accurate.

As a concluding note, 566 techniques and sub-techniques are defined in MITRE ATT&CK [9]. The evaluation was limited to techniques only; logically, sub-technique filings would be similarly categorized as their parent technique. However, further evaluation of all 566 techniques is possible and probably valuable, and future work could concentrate on adding additional techniques and sub-techniques to the mapping. Current findings suggest that there will likely not be a significant deviation from the current findings, but this was not confirmed within the scope of this research. Expanding to more techniques and adding additional incident reporting taxonomies to the assessment are recommended. Finally, we also recommend that any new incident taxonomy be evaluated against the MITRE ATT&CK techniques, as it appears that the model is sound and adequately assesses an incident taxonomy to sufficient tests of accuracy, efficiency, and efficacy controls.

5. Analysis and Interpretation

Analyzing the findings reveals many insights about each of the evaluated incident taxonomies that can assist in modernizing each of the schemas or developing a new schema outright. At a high level, while some taxonomies appeared to perform better in assessment than others regarding potential redundancies in categorization that would lead to mislabeling or analyst confusion, it is notable that all three schemas did experience some redundant labeling, indicating no taxonomy perfectly performed when assigning single categories to each of the assessed techniques. For instance, the Resource Development technique family was not adequately categorized by any of the evaluated taxonomies. An important caveat also worth highlighting is that the duplication of categorization is not specifically an indication of poor category definitions or multiple interpretations but is highly situational depending on the context of how the technique would be notionally applied in an attack path. For instance, "User Execution" is defined by MITRE ATT&CK as an "Execution" technique, resulting in a number of duplicative categorizations across all taxonomies, but this is not only possible due to inefficiencies in said taxonomies. "User Execution" as a technique may involve a user clicking a link in a phishing email, coerced into executing via social engineering or a web application attack, depending on the nature of the attack path used by an adversary. To assist in further analysis, the findings of this taxonomy cross-walk have been sorted into a number of key findings, presented here:

5.1. Overly Specific or Focused Category Conventions

In general, it appears that any incident taxonomy should not attempt to classify an incident by a single event within the wider context of said incident. This is most apparent in the CISA/US-CERT taxonomy, where "Attack Vector" is a single step in the wider attack path. SOCs attempting to categorize incidents later on in the kill chain are forced to categorize the incident as "Unknown" until sufficient timelines are built out wherein relabeling by the correct Attack Vector can occur. In other words, limiting categorizations to a portion of the wider kill chain virtually guarantees that certain actions are "Unknown" depending on where in the attack path the compromise is first detected. For example, should a SOC analyst detect exfiltration of a network share drive to external adversary-controlled infrastructure, there is significant analysis work to be completed before the SOC can confidently

report where the initial access and initial attack even occurred. Considering an incident investigation can take hours, days, and sometimes weeks, this would mean that the incident category remains "Unknown" for an inefficient and unacceptable amount of time.

This is arguably less severe but still as prevalent within the DoD CJCSM 6510.01B construct, as a detected event may not necessarily immediately reveal whether the wider incident involves a "User-Level Compromise" or the notionally more severe "Root Level Compromise." For instance, if a SOC detects "Command & Control" activity emanating from a laptop to adversary infrastructure, it will not be immediately clear whether the compromised account is at the user or administrator level, but reporting time constraints remain. Left to interpretation, some sub-organizations may elect to report as a user-level compromise until it can be confirmed that the root level is achieved. Other sub-organizations may elect to assume root level until disproven. This lack of clarity and accompanying ambiguity is problematic for a top-level organization receiving these reports with limited resources to dedicate to competing incident priorities.

5.2. Duplicating / Conflicting Category Definitions

Likewise, lack of clarity in categorization definitions or overlap in definitions should be avoided, as demonstrated in the vertical clustering apparent in the DoD CJCSM 6510.01B labeling. In addition to the ambiguity caused by a lack of confirmation over a root vs. user-level compromise, the addition of CAT-7 Malicious Logic provides rampant instances of duplicate technically accurate categorizations. While CAT-1 Root Level Compromise and CAT-2 User Level Compromise are relatively self-evident, to reiterate, a CAT-7 Malicious Logic is defined as, *Installation of software designed and/or deployed by adversaries with malicious intentions to gain access to resources or information without the consent or knowledge of the user. This only includes malicious code that does not provide interactive remote control of the compromised IS. Malicious code that has allowed interactive access should be categorized as Category 1 or Category 2 incidents, not Category 7* [6].

While the primary caveat within this definition is the requirement that the identified malicious code does not provide interactive remote control of the system, again, this is a time and detection-dependent detail that a SOC may not have identified at the time of initial alerting. As a result, this adds an additional layer of ambiguity to the DoD model, as made evident in the crosswalk, as nearly all steps post-Initial Access within MITRE ATT&CK could

feasibly be a CAT-1, CAT-2 or a CAT-7 (or none of the above) depending on interpretation and the details available at the time of the defect.

Ambiguity and overlap in definitions exist within the CISA/US-CERT and Verizon taxonomies. For CISA, common attack use cases exist where an attacker may impersonate a legitimate user or service in a phishing email, coercing a user to navigate to a hostile web page. As an attack vector, this use-case is feasibly categorized as "Web" for the hostile link, "Email/Phishing" for the delivery of the link via email, or "Impersonation/Spoofing" since the attacker posed as a legitimate user or service. For Verizon, similarly, in a use-case in which a user receives a malicious email and installs a rootkit, the initial detection may be labeled as "Social Engineering" for the phishing or as "System Intrusion" for the rootkit installation. This finding appears most frequently in the Initial Access and Execution ATT&CK technique families, presumably because all three assessed taxonomies exhibit the most definition ambiguity within this section.

Presumably, variations exist in this portion of the attack path because, logically, this is where most variation in detection occurs. Namely, once an attacker has compromised a system, it's relatively simple to label the incident an "Intrusion," as DoD and Verizon do within their taxonomies. Variation at the Initial Access and Execution levels presupposes that incidents are primarily detected at this portion of the kill chain, which may be idealistic and misguided. As a result, we recommend that CISA/US-CERT relegate the Attack Vector categorization as a secondary field and conceptualize a better mechanism for categorizing incidents, given the breadth of "Unknown" labels that exist across the majority of the cyber kill chain. A potential recommendation for categorization schema may be to utilize the Verizon "System Intrusion" mechanism with a subcategorization that specifically labels the ATT&CK family/technique to granularize further and typify the detected activity.

5.3. User vs. Admin Level Compromise

One noted strength of the taxonomy previously discussed is the explicit differentiation between User-level compromise and Root-level compromise as defined by DoD. This may seem contradictory, as it was highlighted as an ambiguous weakness in prior sections. Still, it bares clarifying that if this level of detail is available, it is indeed a strength to be able to report the privilege level of the compromise. Thus, the ambiguity mentioned above results from timing and circumstance for the individually filed report and

is not necessarily a poor choice to pursue outright. The ability to quickly identify the credential level of a compromise at a glance is arguably useful. For governing organizations like USCYBERCOM, which would receive these incidents, the top-level categorization scheme immediately communicates the priority level, which informs decisions regarding resources, response levels, etc. For internal stakeholders and management, this top-level categorization quickly answers questions surrounding whether the incident has compromised the entire network or if a single system can be quarantined and handled. These are critical differences that exponentially adjust the reactive severity level.

This presents a challenge, considering prior findings regarding adequately selecting a categorization label if the required information is not yet identified. Essentially, in removing the data field from the context of incident taxonomy, the requirement appears to be a need to know the credential level of compromise (user vs. root) as soon as possible. There is little doubt regarding the importance of having this information as soon as it is identified, but communicating it via the incident category is untenable and finding dependent. Thus, we recommend a more generalized labeling scheme, such as Verizon's generic "System Intrusion," with an additional field denoting the level of compromise that can be submitted as soon as this level of detail is known.

5.4. Pre-Attack vs Post-Attack

Another core insight is specific to differentiating between events that constitute "pre-attack adversary activity" and "post-attack adversary activity." MITRE ATT&CK distinguishes this split relatively simply within its "Mitigation" and "Platform" fields across the framework. Each of these fields labels certain Technique Families as "Pre-Attack" and includes "Reconnaissance" and "Resource Development." Predictably, two of the three taxonomies (CISA and Verizon) do not adequately map to any techniques found in these technique families. This is sensible, as one would question why cyber incident categorizations exist for activity that is not specifically regarded as a "cyber incident." This is visible within the assessment tables, notably in Table 1, where "Resource Development" was unable to be mapped due to null results, "Reconnaissance" has differing results for CISA depending on the technique, and Verizon lists all events as "Everything Else."

However, DoD differs from the other two taxonomies in this regard because it has a specific incident category dedicated to "Reconnaissance" (CAT-

6). In practice, this likely has limited success from a detection standpoint, as only "Active Scanning" and similar detectable techniques would be identified. Presumably, it is virtually impossible to identify adversary activity from legitimate activity for the technique "Search Open Websites," for example. Despite this technical shortcoming, DoD finds enough value in categorizing Reconnaissance attempts against its networks and assets that it has established a stand-alone category. Arguably, given a large enough data set coupled with additional inputs such as potential attributions and threat intelligence, this data category is likely valuable as a means of identifying potential future attacks and adversary targeting action. Differentiating this family of techniques from "system intrusion" related events is similarly logical. If an organization is interested in tracking this level of activity, a stand-alone category is likely the best option to do so.

About "Resource Development," from an event detection standpoint, there is likely very little opportunity or value-add in attempting to categorize this technique family. Per MITRE ATT&CK [9], resource development consists of techniques that involve adversaries creating, purchasing, or compromising/stealing resources that can be used to support targeting. Such resources include infrastructure, accounts, or capabilities. These resources can be leveraged by the adversary to aid in other phases of the adversary lifecycle, such as using purchased domains to support Command and Control, email accounts for phishing as a part of Initial Access, or stealing code signing certificates to help with Defense Evasion. In plain terms, this technique family refers to steps an adversary can take to prepare infrastructure set up adjacent accounts for potential impersonation, stage capabilities for later use, etc.

Since the adversary is not specifically attacking or even "touching" the target network during the stage of the attack path, there is very little opportunity to detect these events or if categorization is even appropriate. Indeed, the presence of adversary infrastructure or actions as a potential threat vector aligns very accurately with what "Cyber Threat Intelligence" is intended. NIST 800-150 [16] defines this as "Threat information that has been aggregated, transformed, analyzed, interpreted, or enriched to provide the necessary context for decision-making processes." Information surrounding adversary infrastructure, preferred capabilities, and various machinations occurring before an attack's inception is arguably a firm candidate for what constitutes "Threat Intelligence" and should be regarded as such. As a result, we recommend specifically categorizing any Resource Development activity

via threat intelligence processes and not within any incident reporting process.

5.5. Category Swaps Appear to Come With Challenges

Within the guidance provided for CISA and DoD's taxonomies, special consideration is made for the potential need to update an incident report record, re-categorizing it as appropriate if new information comes to light or if errors are made. This is most evident in CISA's "Unknown" category and its potential to move to any other category and within DoD as the more ambiguous "CAT-8: Investigating" category, which is designed to be the transitive state that a record exists in before a more appropriate category can be named. Within DoD, a CAT-8 can move to any other incident category, or if deemed benign or a false positive, can be moved to "CAT-9: Explained Activity". This approach appears functional, treats the incident record as a living data record, and allows for flexibility in the incident handling process.

That said, unless very robust versioning strategies are being utilized within the database, top-level recategorization (outside of a simple "true" to "false" positive correction) appears messy if organizations wish to ever retrospectively review how incidents are iteratively handled, as the last categorization is the "final say" so to speak. In other words, in taxonomies such as DoD and Verizon, context is lost if an incident record initially starts as a "CAT-7 Malicious Logic" for DoD or "Basic Web Application Attack" for Verizon and eventually moves to a category more severe. It is retrospectively insightful to know if analysts originally identified an event as a certain category and if new information updates the most accurate category for the representative incident. Once again, relegating this specific classification of information to another data field is likely preferred, to the point where a ranked system of categorization is the most preferred mechanism of classifying these various data fields.

5.6. Generalized Labels are More Adaptable

Considering the findings and recommendations thus far, it is becoming increasingly apparent that the more generic the category is, the more adaptable and flexible it is. In truth, following the Execution phase within MITRE ATT&CK, it appears largely dependent on organizational preference regarding the level of detail required for the incident category. Regardless of what additional detail an organization would prefer to provide at the top-level category, the event post-execution is still a "System Intrusion" through the attack

path to the culmination of "Exfiltration" and "Impact," to use Verizon's parlance. As a result, keeping the categorization as simple as possible yields the most adaptability and accuracy throughout the incident handling process, with minimal top-level recategorizations occurring. Granted, redundancy exists within the Verizon schema, particularly with clustering surrounding "Basic Web Application Attacks" and "Social Engineering." However, these appear to exist primarily within the "Initial Access" technique family and subsequently suffer from the same ambiguity here as CISA's schema.

In truth, it appears as if "System Intrusion" is the most accurate category regarding successful attacks. Any other categorizations that characterize the attack vector specifically or other explicit steps within an attack encounter opportunities for ambiguity or overlap in definition interpretation. A more binary approach is preferred, focusing on whether or not the incident occurred, without muddying the categorization scheme with additional extraneous details.

5.7. Key Finding

Overall, in all assessed taxonomies, there appears to be some conflation between "tactical facts" as a function of incident response and "overarching assessments" as a function of threat intelligence. "Overarching Assessments" define the activity in totality, but this requires understanding all steps in the attack path before an accurate label can be applied. "Tactical Facts" refers to a point-in-time identification of a trait within the larger attack path. Since an attack is iterative and multi-step, a static label cannot be applied if the activity detected by an analyst is not within the identified phase of the attack, such as the use-case where "Lateral Movement" detection would result in an "Unknown" label for CISA/US-CERT, as Lateral Movement is further down the kill-chain. Some schema takes this into account, demanding updates to categorization as new information is obtained. It is not immediately clear when a "Basic Web Application Attack" becomes a "System Intrusion" in the Verizon construct, but a "CAT-8 Investigating" in the DoD construct easily moves to another category as the event is investigated. However, if an organization values the collection of metrics on initial detections, a change in category inhibits that ability unless versioning is applied to the record.

Thus, there appears to be inhibitive ambiguity surrounding the intent of the top-level categorization taxonomy that is likely dependent on the organization's priorities. CISA prioritizes the attack vector of an attack, but

this negatively impacts the speed at which reports are communicated unless it is acceptable to label all incidents as "Unknown." Verizon utilizes a generic labeling scheme that yields flexibility for all detected activity (with a few contradictory exceptions). Still, some value appears lost if definitions are too high-level, such as the ability to differentiate between a user-level and root-level compromise. DoD appears to attempt to accomplish both at once, but as a result, contradictions in categorizations exist throughout the taxonomy. Any improvements to existing taxonomies or proposals for new taxonomies should consider these insights during modeling and development.

6. The Proposed Taxonomy for Cyber Incidents

This section specifically seeks opportunities for the practical application of the insights derived from this study. Considering the above, there is potential for a simple top-level taxonomy schema that avoids any potential duplication of incident categorization, minimizes the opportunity to label any top-level incident as "Unknown," and improves on the overarching accuracy and efficiency in declaring and naming an incident with a concise label name. The intent in conceptualizing a new schema is to provide an open and publicly available taxonomy that is effective and efficient for the wider cyber defense community. As mentioned, limited taxonomies were available for study within this work, likely because many organizations use their internal processes with no inherent responsibility to publicize them. We propose UCIT as a potential universal standard for the wider benefit of the community.

6.1. Consolidating Insights from the Evaluation

To reiterate, the primary challenge is developing a new cyber incident taxonomy. There appear to be contradictory requirements in question, both entirely legitimate but seemingly at odds with each other, resulting in ambiguity, inefficiency, and definition overlap. Namely, one use-case treats all incident reports as concluded retrospective timelines of past events, which includes all available technical context. In this regard, an incident report is an "overarching assessment" of the individual events within an attack path that make up the cyber incident. This use case is at odds with incident reporting as a rapid notification system, in which "tactical facts" are reported as they become known. In other words, there is inherent value in having the ability to look back at past cyber incidents and quickly derive an understanding of

what occurred by reading the top-level categorization provided. Likewise, however, there is inherent value in providing a categorization schema to a responding organization to quickly and accurately communicate what they are observing without confusion or uncertainty on the part of any associated stakeholder audience.

Thus, the persistent challenge for all organizations is balancing these competing priorities, considering both are seemingly valid and legitimate needs for the operation. At its core, extremely generalized labeling appears to be the most accurate and pragmatic approach for incident reports as tactical facts. However, these generalizations do not provide rapid detail insights at a glance into how some more contextual incident report category labels achieve, which is to the benefit of the retrospective. Attempts to bridge the gap within the same data field appear to result in filing confusion and indistinct definitions, leading to duplication and conflicting interpretation. Finally, there appears to be no discernible reason why these single data fields require all of this extra contextual information when multi-field sorting and searching is relatively simple to implement and can lead to wider degrees of accuracy.

6.2. Linnaean Taxonomy as a Model

Carl Linnaeus was a Swedish botanist, zoologist, taxonomist, and physician who formalized binomial nomenclature, the modern system of naming organisms, and is considered the father of modern taxonomy. Subsequently, Linnaean Taxonomy [20] is the rank-based classification of organisms or rank-based scientific classification that refers to the mechanism for classifying groups of biological organisms. Current iterations of the taxonomy include the ranks of domain, kingdom, phylum, class, order, family, genus, and species. Groups of a given rank can be aggregated to form a more inclusive group of higher rank, thus creating a taxonomic hierarchy. Simply put, this is a system of categorizing information at varying altitudes of description for the purposes of organizing biological organisms into a ranked hierarchy for clean classifications and sorting of relationships.

Based on the findings of this study, particularly open discussion surrounding the right level of detail in a top-level categorization as well as some of the imperfect ambiguity that comes from attempting to utilize a single data level for multiple purposes, a rank-based taxonomy model similar to that of [20] will potentially address all findings. While biological taxonomy currently has 8 ranks in the hierarchy, an eight-field description is probably infeasible for a

quick one-line synopsis of a cyber incident. However, utilizing a system such as binomial nomenclature or trinomial nomenclature may be more palatable to a stakeholder audience seeking quick insight into the context of a reported cyber event. This would allow for a highly modular taxonomy that can be added depending on the level of detail required by the filing or receiving organization. Additionally, this appears to solve the issue surrounding varying levels of detail adequately. Much like [9] for ranking biological organisms, utilizing such an approach for cyber incidents allows for the top rank of the hierarchy to be universal enough to satisfy adaptability and flexibility needs, with iteratively more detail as the additional ranks in the hierarchy being populated.

However, Linnaean taxonomy is still just a standard taxonomy like any other. The specific usage of Linnaean as a model refers to the specific binomial or trinomial nomenclature in order to fully describe an observed incident succinctly in a "one-liner" for reporting purposes as well as retrospective statistical metrics gathering. The distinction made between standard taxonomy and Linnaean is the specific tendency within Linnaean taxonomies to utilize "Genus-Species" as the common nomenclature for communication. The identified common weakness in the assessed organizational schema was the tendency to attempt to describe an observed cyber incident via a single word or phrase, which resulted in ambiguity, overlap, and potential confusion at the organizational level. A proposed taxonomy might utilize the corresponding "Genus" descriptor as a means of categorizing an event as generally as possible, followed by a "Species" descriptor that more accurately describes the observed event. For example, if the observed event is an attempted "System Compromise," this could refer to any number of offensive techniques, while a "Driveby Compromise" would adequately describe exactly what was first observed. When described binomially, "System Compromise – Driveby Compromise" categorizes the top-level attempt generally, with the added detail of exactly what was observed that would notionally result in the aforementioned end-state.

6.3. The Universal Cyber Incident Taxonomy (UCIT)

Our proposed taxonomy, Universal Cyber Incident Taxonomy (UCIT) (see Table 5), intends to mimic the rank-based hierarchy of the Linnaean model for biology via aggregating data into groupings of increasing rank. UCIT intends to adequately summarize all potential observable events within an adversary's attack path into incrementally detailed hierarchy levels so that

a reported cyber incident can rapidly communicate to stakeholders exactly what it is and what is happening. This is accomplished via a trinomial nomenclature-styled approach to a naming convention, as shown in Table 5. This system relies on a simplistic three-part naming scheme as the top-level categorization framework for any observable event. Namely, the taxonomy relies on an incident Category (or, what is it?), the initial technique family detected (or, what did the analyst first observe that resulted in the decision to report as an incident?), and the currently confirmed level of compromise (User, Root, or Unknown). Special consideration was made to take advantage of a horizontal categorization schema to minimize the potential scope creep via vertical categorizations that were apparent within the various preexisting assessed frameworks.

UCIT utilizes trinomial nomenclature to quickly typify an incident as an adversary's offensive action by simply categorizing it as a "System Intrusion" along with the Technique Family first observed that triggered subsequent filing of the incident. This is based on the results of the cross-walk, in which clustering appeared consistent based on the Technique Family in which the individually assessed Techniques belong assigned without any apparent deviation from the same techniques within the same technique family. This was an unanticipated finding during the testing, as it was assumed that newly defined fields might have to have been developed to account for variations in detected cyber events. However, once it became clear that techniques served as a useful way to categorize an event, it became clear that any new taxonomy should similarly utilize the ATT&CK framework for observed labeling activity. After all, this was the intent of ATT&CK's development. At first, attempts were made to categorize techniques individually. Still, for a summative taxonomy, this quickly proved infeasible, as it would expect reporting SOCs to maintain a database schema with 466 individual techniques. Due to the lack of variation within the higher-level technique family, this was quickly identified as a potential candidate for secondary categorization.

6.3.1. The UCIT Schema Explained

In addition, utilizing MITRE ATT&CK's Technique Family field as a categorization data field serves to quickly identify where in the incident investigation timeline the reporting SOC is analyzing events. This subsequently signals to stakeholders whether the SOC is working "top to bottom" of an attack path (such as if Initial Access was detected), "bottom to top" (such as if Command and Control were detected), or "middle-out" (for instance,

if Lateral Movement was the first detection). This technique rapidly communicates what steps still need to be confirmed. The third field, "Level of Intrusion," borrows from the identified strength of the DoD framework, in which Root or User level is confirmed, but allows for added flexibility in "Unknown" that can be updated later as new information is analyzed. For example, a cyber incident recorded as System Intrusion – Command and Control – User would signal that malware emanating Command and Control traffic out of the network was detected, and stakeholders would quickly derive that all attack-path steps before and after this event are still being investigated.

Both Verizon and DoD have standalone categories for Denial of Service (DoS), which was consolidated along with *Reconnaissance* (from the DoD model) into a category dubbed *Boundary Activity* based on their inherent similarities. Both of these events do not typically involve an accompanying "System Intrusion," so activity that is observed at a network boundary can be categorized appropriately. The Initial Family Detection of "Reconnaissance" or "Impact (Network or Endpoint DoS)" differentiates between the two, with a Level of Intrusion of "N/A" since a User or Root level compromise would lead to an adequate categorization of "System Intrusion" at the top-level. As an example, a cyber incident recorded as Boundary Activity – Reconnaissance – N/A would signal that the reporting SOC has identified that an entity is a vulnerability scanning its network. This would also allow for rapid recategorization to Boundary Activity – Impact – N/A if the scanning was significant enough to knock services temporarily offline for a Denial of Service.

A fourth top-level field of "Report Status" aims to mitigate the top-level recategorization built into all three taxonomies and should be considered a relegated secondary field. "Open/Investigating," "Closed," "Benign," and "Unsuccessful" intend to supplant the recategorization that occurs most apparently in the DoD model with "CAT-3 Unsuccessful Activity", "CAT-8 Investigating" and "CAT-9 Explained Activity" that would complicate the retrospective analysis of past events. It is arguably more valuable to collect past events with what they were initially investigated as only to discover later they were benign or unsuccessful than it would lose that context upon top-level recategorization. For instance, a stakeholder seeking to identify the past year's investigations involving "Lateral Movement" or to use the DoD model, all suspected "CAT-2 User Compromises" would need access to all versioning of the given filed incidents to review this data, as simply pulling

all "CAT-2s" would only collect successfully confirmed compromises and not false positives or unsuccessful attempts, due to the subsequent recategorization. Attempting to pull "CAT-3s" and "CAT-9s" would also collect all other previously categorized incidents other than "CAT-2s" without access to all versions of the database record. Horizontally categorizing this field is arguably a better data practice and avoids this complication.

"Non-Compliance Activity" was added as a means to account for DoD's "CAT-5 Non-Compliance Activity" and Verizon's "Miscellaneous Errors" and some use-cases of "Privilege Misuse." However, these categories are arguably out of scope for this study primarily because they pertain to individual and organizational policy violations or misconfigurations. This study concentrated entirely and explicitly on accurately categorizing offensive cyber actions. It is recommended that if this taxonomy is adopted by any organization, the secondary field encapsulating MITRE ATT&CK technique families be customized for subcategories of individual policy violations such as "PII spillage," "Poor Security Practice," etc. Of note, it is debatable as to whether a Security Operations Center dedicated to detecting hostile, offensive cyber activity should involve itself in policy violations at all. This is likely organizationally dependent. This category is offered to complete the UCIT taxonomy and can be applied or disregarded according to individual organization needs.

Additionally, the performance of CISA/US-CERT's "Attack Vector" based taxonomy was disregarded and not included in the UCIT construct. While the initial Attack Vector is arguably a valuable data field to include in an incident report's required fields, it was deemed infeasible to include in any top-level categorization schema for the reasons outlined in Section V. Notionally, as a SOC completes its investigation into an identified cyber incident, each step of the attack-path will be labeled and reported appropriately. Still, for the purposes of the initial notification, this information may not be available and thus should not be a part of the top-level categorization schema.

6.3.2. Recognizance of "Point-in-Time" Status

While not included in the UCIT table example, it was later identified that while a trinomial naming convention would accurately describe an incident as it was initially reported, the taxonomy lacked a mechanism for accurate documentation once an incident investigation was completed and closed out. This was initially deemed out of scope, as without real incident record examples to utilize, any assessment of cyber incidents in totality would be pure

conjecture at best. This could be potentially rectified by organizations constantly recategorizing their incident records as new techniques were observed, but this is simply infeasible and increases confusion, particularly in regard to how an incident is communicated and reported, considering it would be regularly changing names. However, this weakness could be rectified by requiring the additional field of "Final Observed Sub-Technique" to illustrate the incident's attack path in totality. This would increase the accuracy of the record in totality as well as provide valuable statistical metrics regarding trends and averages in observed attack paths and levels of attacker success. For example, in a notional scenario involving a Drive-by download of malware resulting in a user-level compromise and deletion of a file system, a "System Intrusion – Driveby Compromise" incident record would be used to document investigative steps and would be closed out with the additional field of "Data Destruction" as a sub-technique to illustrate the end-state of the cyber incident.

Finally, UCIT was subjected to the same cross-walk evaluation as the three assessed taxonomies, and results were split between Table 6 and Table 7. Results were precise as expected, which is unsurprising considering the taxonomy uses MITRE ATT&CK as its primary sorting mechanism. Ambiguity does exist in the "Impact" technique family, predictably, as this event is situationally dependent. For instance, a DoS can occur as a "Boundary Action" and after a "System Intrusion" has occurred. However, these top-level categorizations adequately differentiate between these two use cases and should result in minimal confusion from a stakeholder audience.

Table 2: Persistence, Privilege Escalation, and Defense Evasion.

		Persistence						Privilege Escalation						Defense Evasion					
		Scheduled Task/Job	Office Application Startup	Account Manipulation	Create Account	Browser Extensions	Implant Internal Image	Access Token Manipulation	Abuse Elevation Control Mechanism	Exploitation for Privilege Escalation	Escape to Host	Event Triggered Execution	Component Object Model Hijacking	Obfuscated Files or Information	Impair Defenses	Masquerading	Direct Volume Access	Rootkit	Indicator Removal on Host
CISA/US-CERT	Unknown	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	Attrition																		
	Web																		
	Email/Phishing																		
	External/Removable Media																		
	Impersonation/Spoofing																		
	Improper Usage																		
	Loss or Theft of Equipment																		
	Other																		
DOD CJCSM 6510.01B	CAT-1 Root Level Intrusion	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	CAT-2 User Level Intrusion	X	X	X	X	X	X				X	X	X	X	X	X	X	X	X
	CAT-3 Unsuccessful Activity																		
	CAT-4 Denial of Service																		
	CAT-5 Non-Compliance Activity																		
	CAT-6 Reconnaissance																		
	CAT-7 Malicious Logic	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	CAT-8 Investigating																		
	CAT-9 Explained Activity																		
Verizon DBIR	Basic Web Application Attacks																		
	Denial of Service																		
	Lost and Stolen Assets																		
	Miscellaneous Errors																		
	Privilege Misuse																		
	Social Engineering																		
	System Intrusion	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	Everything Else																		

Table 3: Credential Access, Discovery, and Lateral Movement.

		Credential Access						Discovery						Lateral Movement					
		Unsecured Credentials	Brute Force	Forge Web Credentials	OS Credential Dumping	Two-Factor Authentication Interception	Steal Application Access Token	Network Share Discovery	Process Discovery	Query Registry	System Service Discovery	Application Window Discovery	Remote System Discovery	Lateral Tool Transfer	Exploitation of Remote Services	Taint Shared Content	Remote Services	Internal Spearphishing	Remote Service Session Hijacking
CISA/US-CERT	Unknown				X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	Attrition		X																
	Web			X															
	Email/Phishing																	X	
	External/Removable Media																		
	Impersonation/Spoofing	X																	
	Improper Usage	X																	
	Loss or Theft of Equipment																		
	Other																		
DOD CJCSM 6510.01B	CAT-1 Root Level Intrusion	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	CAT-2 User Level Intrusion	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	CAT-3 Unsuccessful Activity																		
	CAT-4 Denial of Service																		
	CAT-5 Non-Compliance Activity																		
	CAT-6 Reconnaissance																		
	CAT-7 Malicious Logic		X	X	X	X	X	X		X	X	X	X	X	X	X	X	X	X
	CAT-8 Investigating																		
	CAT-9 Explained Activity																		
Verizon DBIR	Basic Web Application Attacks																		
	Denial of Service																		
	Lost and Stolen Assets																		
	Miscellaneous Errors																		
	Privilege Misuse	X																	
	Social Engineering																		X
	System Intrusion	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
	Everything Else																		

Table 4: Normalization in Remaining MITRE ATT&CK Families.

		Collection						Command and Control							Exfiltration							Impact						
		Data from Local System	Screen Capture	Archive Collected Data	Data from Network Shared Drive	Email Collection	Browser Session Hijacking	Application Layer Protocol	Non-Standard Port	Ingress Tool Transfer	Data Obfuscation	Fallback Channels	Proxy		Exfiltration Over C2 Channel	Exfiltration Over Web Service	Exfiltration over Physical Medium	Data Transfer Size Limits	Exfiltration Over Alternative Protocol	Transfer Data to Cloud Account		Data Destruction	Data Manipulation	Defacement	Service Stop	System Shutdown/Reboot	Disk Wipe	
CISA / US-CERT	Unknown	X	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X		X	X	X	X	X	X	
	Attrition																					X						
	Web																							X				
	Email/Phishing																											
	External/Removable Media																											
	Impersonation/Spoofing																											
	Improper Usage																											
	Loss or Theft of Equipment																											
	Other																											
DOD CJCSM 6510.01B	CAT-1 Root Level Intrusion	X	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X		X	X	X	X	X	X	
	CAT-2 User Level Intrusion	X	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X		X	X	X	X	X	X	
	CAT-3 Unsuccessful Activity																											
	CAT-4 Denial of Service																											
	CAT-5 Non-Compliance Activity																											
	CAT-6 Reconnaissance																											
	CAT-7 Malicious Logic	X	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X		X	X	X	X	X		
	CAT-8 Investigating																											
	CAT-9 Explained Activity																											
Verizon DBIR	Basic Web Application Attacks																					X	X	X				
	Denial of Service																											
	Lost and Stolen Assets																											
	Miscellaneous Errors																											
	Privilege Misuse																											
	Social Engineering																											
	System Intrusion	X	X	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X		X	X	X	X	X	X	
	Everything Else																											

Table 5: Universal cyber incident taxonomy (UCIT).

Category	Initial Family Detection	Level of Intrusion	Report Status
Boundary Activity	Reconnaissance	N/A	Open / Investigating
	Impact (Network or Endpoint Denial of Service)		
System Intrusion	Initial Access	Root	Closed
	Execution		
	Persistence		
	Privilege Escalation		
	Defense Evasion	User	
	Credential Access		
	Discovery		Benign
	Lateral Movement		
	Collection	Unknown	
	Command and Control		
	Exfiltration		Unsuccessful
	Impact		
Non-Compliance Activity			

Table 6: UCIT crosswalk against MITRE ATT&CK (reconnaissance to defense evasion).

		MITRE ATT&CK Technique (Reconnaissance to Defense Evasion)														
			Reconnaissance				Resource Development				Initial Access					
			Active Scanning	Search Open Websites/Domains	Gather Victim Identity Information	Gather Victim Org Information	Search Victim-Owned Websites	Search Open Technical Databases	Compromise Infrastructure	Establish Accounts	Stage Capabilities	Acquire Infrastructure	Compromise Accounts	Develop Capabilities	Drive-by Compromise	Exploit Public-Facing Application
Boundary Activity	Reconnaissance		X	X	X	X	X	X								
	Impact (Network or Endpoint Denial of Service)														X	X
System Intrusion	Initial Access														X	X
	Execution															
	Persistence															
	Privilege Escalation															
	Defense Evasion															
	Credential Access															
	Discovery															
	Lateral Movement															
	Collection															
	Command and Control															
	Exfiltration															
	Impact															
Non-Compliance Activity																

		MITRE ATT&CK Technique (Reconnaissance to Defense Evasion)														
			Initial Access		Execution				Persistence						Privilege Escalation						Defense Evasion					
			Hardware Additions	Trusted Relationship	Supply Chain Compromise	Phishing	User Execution	Command and Scripting Interpreter	Exploitation for Client Execution	Windows Management Instrumentation	Native API	System Services	Scheduled Task/Job	Office Application Startup	Account Manipulation	Create Account	Browser Extensions	Implant Internal Image	Access Token Manipulation	Abuse Elevation Control Mechanism	Exploitation for Privilege Escalation	Escape to Host	Event Triggered Execution	Component Object Model Hijacking	Obfuscated Files or Information	Impair Defenses	Masquerading	Direct Volume Access	Rootkit	Indicator Removal on Host
Boundary Activity	Reconnaissance																											
	Impact (Network or Endpoint Denial of Service)		X	X																								
System Intrusion	Initial Access		X	X																								
	Execution				X	X	X	X	X	X	X	X																
	Persistence												X	X	X	X	X	X										
	Privilege Escalation																		X	X	X	X	X	X				
	Defense Evasion																		X	X	X	X	X	X	X	X	X	X
	Credential Access																											
	Discovery																											
	Lateral Movement																											
	Collection																											
	Command and Control																											
	Exfiltration																											
	Impact																											
Non-Compliance Activity																												

Table 7: UCIT crosswalk against MITRE ATT&CK (credential access to impact).

		MITRE ATT&CK Technique (Credential Access to Impact)					
		Credential Access					
		Unsecured Credentials	Brute Force	Forge Web Credentials	OS Credential Dumping	Two-Factor Authentication Interception	Steal Application Access Token
Boundary Activity	Reconnaissance						
	Impact (Network or Endpoint Denial of Service)						
System Intrusion	Initial Access						
	Execution						
	Persistence						
	Privilege Escalation						
	Defense Evasion						
	Credential Access	X	X	X	X	X	X
	Discovery						
	Lateral Movement						
	Collection						
	Command and Control						
	Exfiltration						
	Impact						
Non-Compliance Activity							

		Discovery					
		Network Share Discovery	Process Discovery	Query Registry	System Service Discovery	Application Window Discovery	Remote System Discovery
Boundary Activity	Reconnaissance						
	Impact (Network or Endpoint Denial of Service)						
System Intrusion	Initial Access						
	Execution						
	Persistence						
	Privilege Escalation						
	Defense Evasion						
	Credential Access						
	Discovery	X	X	X	X	X	X
	Lateral Movement						
	Collection						
	Command and Control						
	Exfiltration						
	Impact						
Non-Compliance Activity							

		Lateral Movement					
		Lateral Tool Transfer	Exploitation of Remote Services	Taint Shared Content	Remote Services	Internal Spearphishing	Remote Service Session Hijacking
Boundary Activity	Reconnaissance						
	Impact (Network or Endpoint Denial of Service)						
System Intrusion	Initial Access						
	Execution						
	Persistence						
	Privilege Escalation						
	Defense Evasion						
	Credential Access						
	Discovery						
	Lateral Movement	X	X	X	X	X	X
	Collection						
	Command and Control						
	Exfiltration						
	Impact						
Non-Compliance Activity							

		Collection					
		Data from Local System	Screen Capture	Archive Collected Data	Data from Network Shared Drive	Email Collection	Browser Session Hijacking
Boundary Activity	Reconnaissance						
	Impact (Network or Endpoint Denial of Service)						
System Intrusion	Initial Access						
	Execution						
	Persistence						
	Privilege Escalation						
	Defense Evasion						
	Credential Access						
	Discovery						
	Lateral Movement						
	Collection	X	X	X	X	X	X
	Command and Control						
	Exfiltration						
	Impact						
Non-Compliance Activity							

		Command and Control					
		Application Layer Protocol	Non-Standard Port	Ingress Tool Transfer	Data Obfuscation	Fallback Channels	Proxy
Boundary Activity	Reconnaissance						
	Impact (Network or Endpoint Denial of Service)						
System Intrusion	Initial Access						
	Execution						
	Persistence						
	Privilege Escalation						
	Defense Evasion						
	Credential Access						
	Discovery						
	Lateral Movement						
	Collection						
	Command and Control	X	X	X	X	X	X
	Exfiltration						
	Impact						
Non-Compliance Activity							

		Exfiltration					
		Exfiltration Over C2 Channel	Exfiltration Over Web Service	Exfiltration over Physical Medium	Data Transfer Size Limits	Exfiltration Over Alternative Protocol	Transfer Data to Cloud Account
Boundary Activity	Reconnaissance						
	Impact (Network or Endpoint Denial of Service)						
System Intrusion	Initial Access						
	Execution						
	Persistence						
	Privilege Escalation						
	Defense Evasion						
	Credential Access						
	Discovery						
	Lateral Movement						
	Collection						
	Command and Control						
	Exfiltration	X	X	X	X	X	X
	Impact						
Non-Compliance Activity							

		Impact					
		Data Destruction	Data Manipulation	Defacement	Service Stop	System Shutdown/Reboot	Disk Wipe
Boundary Activity	Reconnaissance						
	Impact (Network or Endpoint Denial of Service)				X	X	
System Intrusion	Initial Access						
	Execution						
	Persistence						
	Privilege Escalation						
	Defense Evasion						
	Credential Access						
	Discovery						
	Lateral Movement						
	Collection						
	Command and Control						
	Exfiltration						
	Impact	X	X	X	X	X	X
Non-Compliance Activity							

7. Conclusion and Outlook

We focused on identifying a consistent and accurate incident categorization taxonomy that can be utilized by any organization without substantial customization for unique organizational requirements, as appears to be the case in the current cyber defense ecosystem. As cybersecurity defensive operations have grown into relatively mature business processes, it appears as though the mechanisms for sorting cyber incident data have been extremely fragmented and organizationally specific. While there is no escape from policy governance or legislative requirements, the discrepancies between these three large organizations, both public and private, suggest the opportunity to provide data-driven recommendations for a categorization taxonomy that can be utilized uniformly across both sectors. Currently, these categorizations become the preferred nomenclature for how incidents are defined throughout the organization's incident response process. Still, they are stove-piped within an organization and not necessarily translatable to others. Considering this customization and variance, an opportunity exists to understand differences in how cyber incidents are perceived and accurately labeled.

We evaluated three publicly available categorization taxonomies to understand these differences and assess best practices, potential gaps, and any other challenges that result from self-defined incident categories. Such as DoD, CISA, and Verizon are feasible examples of mature categorization schema. The MITRE ATT&CK framework also offered a robust lexicon of potential cyber events that was used as a data set to assess this schema. The study cross-walked this categorization schema against multiple techniques as defined by MITRE ATT&CK to understand how accurate each schema is, what strengths each schema has, and what gaps can result when various techniques are exercised. Additionally, the study sought to understand similar strengths and weaknesses in categorization workflows, such as if organizations are tasked with changing categorization as new information is presented or circumstances change. Findings suggested consistent patterning of categorizations across all three organizations that suggested either consistent best practices or consistent room for recommendations.

Finally, the study culminated with the development and proposal of the UCIT, which leveraged the study's findings as a potential offering that is publicly available and universally usable by any cybersecurity organization or SOC. Consolidating multiple categories as much as possible and utilizing a system based on Linnaean taxonomy for biological organisms, UCIT offers

a concise and accurate taxonomy that circumvents much of the ambiguity identified in the assessed schema during evaluations. UCIT is designed to be modular, allowing additional fields to be added as required without substantially sacrificing brevity and accuracy. In short, the key contribution of this study is to advocate for optimizing simplicity in an incident reporting schema to eliminate any aforementioned ambiguity. Using stacked categories with incremental detail and relying on the industry standard MITRE ATT&CK framework for identifying the initially detected event, the overarching categorization of the incident in question is more accurately identified and communicated. Some of the future tasks would be the following.

- Additional testing, particularly in a live environment, is required to further strengthen its proof of concept. A significant number of MITRE ATT&CK techniques remain unevaluated, and all four cyber incident taxonomies would benefit from additional evaluation along these lines. Further peer review for UCIT as a means to identify potential contradictions, logic flaws, and process breakdowns would also further strengthen the taxonomy's viability as a potentially usable data schema.

- It is highly likely that many organizations similarly mandate required data fields within their initial incident reporting, as CJCSM 6510.01B [6] outlines within its guidance. If there are data strategies that can enable top-level reporting that also includes these critical fields, that is also an opportunity for future research.

- While "Noncompliance" and associated policy violation activity was deemed out of scope, it raises the question of whether this activity should be within the scope of a SOC, considering the presumed workload that comes with the detection of cyber activity. While it is sensible to house "Noncompliance" associated detection within a SOC construct due to its access to the supporting data, it may be worthwhile to investigate if a standalone or separate entity should exist within an organization to singularly dedicate itself to detecting policy violations. This would notionally free up a SOC to dedicate itself fully to detecting adversarial activity.

Appendix A.

Table A.8: CJCSM 6510.01B Incident/Event Definitions.

Category	Definition
CAT-1 Root Level Intrusion	Unauthorized privileged access to an IS. Privileged access, often referred to as administrative or root access, provides unrestricted access to the IS. This category includes unauthorized access to information or unauthorized access to account credentials that could be used to perform administrative functions (e.g., domain administrator). If the IS is compromised with malicious code that provides interactive remote control, it will be reported in this category.
CAT-2 User Level Intrusion	Unauthorized non-privileged access to an IS. Non-privileged access, often referred to as user-level access, provides restricted access to the IS based on the privileges granted to the user. This includes unauthorized access to information or unauthorized access to account credentials that could be used to perform users' functions such as accessing Web applications, Web portals, or other similar information resources. If the IS is compromised with malicious code that provides interactive remote control, it will be reported in this category.
CAT-3 Unsuccessful Activity	Deliberate attempts to gain unauthorized access to an IS that is defeated by normal defensive mechanisms. The attacker fails to gain access to the IS (i.e., the attacker attempts valid or potentially valid username and password combinations), and the activity cannot be characterized as exploratory scanning. Reporting these events is critical for the gathering of useful effects-based metrics for commanders.
CAT-4 Denial of Service	Activity that denies degrades or disrupts normal functionality of an IS or DoD information network.
CAT-5 Non-Compliance Activity	Activity that potentially exposes ISs to increased risk as a result of the action or inaction of authorized users. This includes administrative and user actions such as failure to apply security patches, connections across security domains, installation of vulnerable applications, and other breaches of existing DoD policy.
CAT-6 Reconnaissance	Activity seeks to gather information to characterize ISs, applications, DoD information networks, and users that may be useful in formulating an attack. This includes activities such as mapping DoD information networks, IS devices and applications, interconnectivity, and their users or reporting structure. This activity does not directly result in a compromise.
CAT-7 Malicious Logic	Installation of software designed and/or deployed by adversaries with malicious intentions to gain access to resources or information without the consent or knowledge of the user. This only includes malicious code that does not provide interactive remote control of the compromised IS. Malicious code that has allowed interactive access should be categorized as Category 1 or Category 2 incidents, not Category 7.
CAT-8 Investigating	Events that are potentially malicious or anomalous activity deemed suspicious and warranted or are undergoing further review. No event will be closed out as a Category 8. Category 8 will be recategorized to appropriate Category 1-7 or 9 before closure.
CAT-9 Explained Activity	Suspicious events that, after further investigation, are determined to be non-malicious activity and do not fit the criteria for any other categories. This includes events such as IS malfunctions and false alarms. When reporting these events, the reason for which they cannot be otherwise categorized must be specified.

Table A.9: Verizon DBIR incident categories/definitions.

Category	Definition
Basic Web Application Attacks	These attacks are against a Web application, and after initial compromise, they do not have many additional Actions. It is the "get in, get the data, and get out" pattern.
Denial of Service	Attacks intended to compromise the availability of networks and systems. This includes both network and application layer attacks.
Lost and Stolen Assets	Incidents where an information asset went missing, whether through misplacement or malice.
Miscellaneous Errors	Incidents where unintentional actions directly compromised a security attribute of an information asset. This does not include lost devices, which are grouped with theft instead.
Privilege Misuse	Incidents predominantly driven by unapproved or malicious use of legitimate privileges.
Social Engineering	A psychological compromise of a person that alters their behavior into taking action or breaching confidentiality.
System Intrusion	Complex attacks that leverage malware and/or hacking to achieve their objectives, including deploying Ransomware.
Everything Else	This "pattern" is not a pattern at all. Instead, it covers all incidents that don't fit within the orderly confines of the other patterns. Like that container where you keep all the cables for electronics you don't own anymore: Just in case.

Table A.10: CISA/US-CERT attack vector categories/definitions.

Category	Definition
Unknown	Cause of attack is unidentified.
Attrition	An attack that employs brute force methods to compromise, degrade, or destroy systems, networks, or services.
Web	An attack executed from a website or web-based application.
Email/Phishing	An attack executed via an email message or attachment.
External/Removable Media	An attack executed from removable media or a peripheral device.
Impersonation/Spoofing	An attack involving the replacement of legitimate content/services with a malicious substitute.
Improper Usage	Any incident resulting from violating an organization's acceptable usage policies by an authorized user, excluding the above categories.
Loss or Theft of Equipment	The loss or theft of a computing device or media used by the organization.
Other	An attack method does not fit into any other vector.

Table A.11: MITRE ATT&CK Glossary.

ATT&CK Tactic	**Definition**
Reconnaissance	The adversary is trying to gather the information they can use to plan future operations. Reconnaissance consists of techniques that involve adversaries actively or passively gathering information that can be used to support targeting. Such information may include details of the victim's organization, infrastructure, or staff/personnel. This information can be leveraged by the adversary to aid in other phases of the adversary lifecycle, such as using gathered information to plan and execute Initial Access to the scope and prioritize post-compromise objectives or to drive and lead further Reconnaissance efforts.
Resource Development	The adversary is trying to establish resources they can use to support operations. Resource Development consists of techniques that involve adversaries creating, purchasing, or compromising/stealing resources that can be used to support targeting. Such resources include infrastructure, accounts, or capabilities. These resources can be leveraged by the adversary to aid in other phases of the adversary lifecycle, such as using purchased domains to support Command and Control, email accounts for phishing as a part of Initial Access, or stealing code signing certificates to help with Defense Evasion.
Initial Access	The adversary is trying to get into your network. Initial Access consists of techniques that use various entry vectors to gain their initial foothold within a network. Techniques used to gain a foothold include targeted spearphishing and exploiting weaknesses on public-facing web servers. Footholds gained through initial access may allow for continued access, like valid accounts and use of external remote services, or maybe limited use due to changing passwords.
Execution	The adversary is trying to run malicious code. Execution consists of techniques that result in adversary-controlled code running on a local or remote system. Techniques that run malicious code are often paired with techniques from all other tactics to achieve broader goals, like exploring a network or stealing data. For example, an adversary might use a remote access tool to run a PowerShell script that does Remote System Discovery.
Persistence	The adversary is trying to maintain their foothold. Persistence consists of techniques that adversaries use to keep access to systems across restarts, changed credentials, and other interruptions that could cut off their access. Techniques used for persistence include any access, action, or configuration changes that let them maintain their foothold on systems, such as replacing or hijacking legitimate code or adding startup code.

Privilege Escalation	The adversary is trying to gain higher-level permissions. Privilege Escalation consists of techniques adversaries use to gain higher-level permissions on a system or network. Adversaries can often enter and explore a network with unprivileged access but require elevated permissions to follow through on their objectives. Common approaches are to take advantage of system weaknesses, misconfigurations, and vulnerabilities.
Defense Evasion	The adversary is trying to avoid being detected. Defense Evasion consists of techniques adversaries use to avoid detection throughout their compromise. Techniques used for defense evasion include uninstalling/disabling security software or obfuscating/encrypting data and scripts. Adversaries also leverage and abuse trusted processes to hide and masquerade their malware. Other tactics' techniques are cross-listed here when those techniques include the added benefit of subverting defenses.
Credential Access	The adversary is trying to steal account names and passwords. Credential Access consists of techniques for stealing credentials like account names and passwords. Techniques used to get credentials to include keylogging or credential dumping. Using legitimate credentials can give adversaries access to systems, make them harder to detect, and provide the opportunity to create more accounts to help achieve their goals.
Discovery	The adversary is trying to figure out your environment. Discovery consists of techniques an adversary may use to gain knowledge about the system and internal network. These techniques help adversaries observe the environment and orient themselves before deciding how to act. They also allow adversaries to explore what they can control and what's around their entry point to discover how it could benefit their current objective. Native operating system tools are often used toward this post-compromise information-gathering objective.
Lateral Movement	The adversary is trying to move through your environment. Lateral Movement consists of techniques that adversaries use to enter and control remote systems on a network. Following through on their primary objective often requires exploring the network to find their target and subsequently gaining access to it. Reaching their objective often involves pivoting through multiple systems and accounts to gain. Adversaries might install their remote access tools to accomplish Lateral Movement or use legitimate credentials with native network and operating system tools, which may be stealthier.

Collection	The adversary is trying to gather data of interest to their goal. The collection consists of techniques adversaries may use to gather information and the sources of information collected that are relevant to following through on the adversary's objectives. Frequently, the next goal after collecting data is to steal (exfiltrate) the data. Common target sources include various drive types, browsers, audio, video, and email. Common collection methods include capturing screenshots and keyboard input.
Command and Control	The adversary is trying to communicate with compromised systems to control them. Command and Control consist of techniques that adversaries may use to communicate with systems under their control within a victim network. Adversaries commonly attempt to mimic normal, expected traffic to avoid detection. There are many ways an adversary can establish command and control with various levels of stealth depending on the victim's network structure and defenses.
Exfiltration	The adversary is trying to steal data. Exfiltration consists of techniques adversaries may use to steal data from your network. Once they've collected data, adversaries often package it to avoid detection while removing it. This can include compression and encryption. Techniques for getting data out of a target network typically include transferring it over their command and control channel or an alternate channel. They may also include putting size limits on the transmission.
Impact	The adversary is trying to manipulate, interrupt, or destroy your systems and data. Impact consists of techniques that adversaries use to disrupt availability or compromise integrity by manipulating business and operational processes. Techniques used for impact can include destroying or tampering with data. In some cases, business processes can look fine but may have been altered to benefit the adversaries' goals. These techniques might be used by adversaries to follow through on their end goal or to provide cover for a confidentiality breach.

References

[1] K. Julisch, Understanding and overcoming cyber security anti-patterns, Computer Networks 57 (10) (2013) 2206–2211.

[2] P. Cichonski, T. Millar, T. Grance, K. Scarfone, NIST special publication 800-61 rev 2: computer security incident handling guide, National Institute of Standards and Technology, Gaithersburg MD (2012).

[3] B. Caskurlu, A. Gehani, C. C. Bilgin, K. Subramani, Analytical models for risk-based intrusion response, Computer Networks 57 (10) (2013) 2181–2192.

[4] A. Ahmad, S. B. Maynard, K. C. Desouza, J. Kotsias, M. T. Whitty, R. L. Baskerville, How can organizations develop situation awareness for incident response: A case study of management practice, Computers & Security 101 (2021) 102122.

[5] M. Grønberg, An ontology for cyber threat intelligence, Master's thesis (2019).

[6] U.S. Department of Defense, Chairman of the joint chiefs of staff manual, cyber incident handling program: CJCSM 6510.01B, 2012.

[7] CISA/US-CERT, Federal incident notification guidelines, `[Online]. Available: https://www.cisa.gov/uscert/incident-notification-guidelines/attack-vectors`, [Accessed: Dec. 6, 2022] (2021).

[8] Verizon, Incident classification patterns, `[Online]. Available: https://www.verizon.com/business/resources/reports/dbir/2021/incident-classification-patterns/`, [Accessed: Dec. 6, 2022] (2021).

[9] B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, C. B. Thomas, Mitre att&ck: Design and philosophy, in: Technical report, The MITRE Corporation, 2018.

[10] P. Lif, S. Varga, M. Wedlin, D. Lindahl, M. Persson, Evaluation of information elements in a cyber incident report, in: Proc. IEEE European Symposium on Security and Privacy Workshops (EuroS&PW), 2020, pp. 17–26. doi:10.1109/EuroSPW51379.2020.00012.

[11] F. Menges, G. Pernul, A comparative analysis of incident reporting formats, Computers & Security 73 (2018) 87–101.

[12] P. Lif, T. Sommestad, D. Granasen, Development and evaluation of information elements for simplified cyber-incident reports, in: Proc. IEEE International Conference On Cyber Situational Awareness, Data Analytics And Assessment (Cyber SA), 2018, pp. 1–10. doi:10.1109/CyberSA.2018.8551402.

[13] S. J. Zaccaro, A. K. Hargrove, T. R. Chen, K. M. Repchick, T. McCausland, A comprehensive multilevel taxonomy of cyber security incident response performance, in: Psychosocial dynamics of cyber security, Routledge, 2016, pp. 43–85.

[14] J. Yuill, F. Wu, J. Settle, F. Gong, R. Forno, M. Huang, J. Asbery, Intrusion-detection for incident-response, using a military battlefield-intelligence process, Computer Networks 34 (4) (2000) 671–697.

[15] B. Zhu, A. Joseph, S. Sastry, A taxonomy of cyber attacks on scada systems, in: 2011 International conference on internet of things and 4th international conference on cyber, physical and social computing, IEEE, 2011, pp. 380–388.

[16] M. D. Ibrishimova, Cyber incident classification: issues and challenges, in: International Conference on P2P, Parallel, Grid, Cloud and Internet Computing, Springer, 2018, pp. 469–477.

[17] S. Cho, I. Han, H. Jeong, J. Kim, S. Koo, H. Oh, M. Park, Cyber kill chain based threat taxonomy and its application on cyber common operational picture, in: 2018 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (Cyber SA), IEEE, 2018, pp. 1–8.

[18] K. W. G. Ang, A case study for cyber incident report in industrial control systems, Ph.D. thesis, Massachusetts Institute of Technology (2022).

[19] B. Al-Sada, A. Sadighian, G. Oligeri, Analysis and characterization of cyber threats leveraging the mitre att&ck database, IEEE Access (2023).

[20] P. Fara, Sex, botany and empire: the story of Carl Linnaeus and Joseph Banks, Icon Books, 2004.