



SANS Institute

Information Security Reading Room

Evaluation of Comprehensive Taxonomies for Information Technology Threats

Steven Launius

Copyright SANS Institute 2020. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Evaluation of Comprehensive Taxonomies for Information Technology Threats

GIAC (GCIA) Gold Certification and RES 5500

Author: Steven M. Launius, stevelaunius@gmail.com

Advisor: *Randy Marchany*

Accepted: *March 1, 2018*

Abstract

Categorization of all information technology threats can improve communication of risk for an organization's decision-makers who must determine the investment strategy of security controls. While there are several comprehensive taxonomies for grouping threats, there is an opportunity to establish the foundational terminology and perspective for communicating threats across the organization. This is important because confusion about information technology threats pose a direct risk of damaging an organization's operational longevity. In order for leadership to allocate security resources to counteract prevalent threats in a timely manner, they must understand those threats quickly. A study that investigates categorization techniques of information technology threats to non-technical decision-makers through a qualitative review of grouping methods for published threat taxonomies could remedy the situation.

1. Introduction

A modern organization's operations depend on information technology (IT). Ubiquitous adoption of IT due to technological advancements creates both efficiencies and vulnerabilities in an organization's operations. Physical threats to IT infrastructure from both human and environmental sources have remained mostly consistent over time. The continuous development of IT systems for exchanging, processing, and storing information introduces many weaknesses. Criminals, activists, nation-states, and other adversaries are increasingly successful at attacking these systems to accomplish their objectives. Many organizations are adopting Cyber Threat Intelligence (CTI) to address the increase in adversarial cyber threats. Since the primary use of CTI is the sharing of an adversary's activities, several taxonomies and ontologies exist for maintaining a common lexicon within and between organizations.

However, in addition to nefarious humans, sources of IT threats may also be accidental, environmental, political, or economical. Leadership must evaluate risk to IT by assessing the likelihood of threat events from all of these sources and their impact on the organization. Risk management professionals from the information security community have published comprehensive taxonomies for grouping threats events. Each taxonomy presents a hierarchy of discrete threat event groups with succeeding levels providing terms with more detail. Categorization and definitions of terms for threat events support communication with decision makers who must select a course of action to counter a threat.

A threat taxonomy can improve communication in two ways. First, language barriers between professionals with different expertise can be broken down into clear definitions for IT threats. As mass media quickly spreads news of IT failures, like cyberattacks or data breaches, a foundation of terms can help decision-makers understand the active threats. Second, an ordered taxonomy structure of the entire IT threat landscape enables analysis and assessment at various granularities. Comparing the risk of high-level threat categories can empower leadership to make the right decisions to protect their organization.

Steve Launius

2. Communicating Threat

2.1. Threat Language

Language is an intricate cognitive process requiring an agreement of standard definitions for effective communication. While the English language has broadly held standards, there are many deviations that can present communication problems. In particular, slang differences occur at many levels:

- National: Americans live in *apartments*, while Brits live in *flats*.
- Regional: Soda, pop, coke, and soft drink are all terms for a sweetened carbonated beverage.
- Local: In Texas, a nag is called a *worrit*.
- Professional: In the health profession, a *virus* is a microorganism that infects living cells to live and reproduce itself and causes human illness (Definition of Virus, 2018). In the IT profession, a *virus* is a hidden, self-replicating section of computer software, usually malicious logic, propagating by infection of another program (Glossary of Security Terms, 2018).

Adhering to standard definitions for threat terms can improve comprehension of the dialog between echelons in any organization. There is no authoritative source for IT threat terms, but there are several glossaries or lexicons of security terms published by a variety of governing bodies. The United States (US) government alone has many sources including:

- Department of Defense (DOD) - Dictionary of Military and Associated Terms,
- Department of Homeland Security (DHS) - Risk Lexicon,
- National Institute of Standards and Technology (NIST) - Glossary of Key Information Security Terms,
- Committee on National Security Systems (CNSS) - Glossary, and
- National Initiative for Cybersecurity Careers and Studies (NICCS) - A Glossary of Common Cybersecurity Terminology.

Many information security organizations also maintain security term definitions:

- SysAdmin, Audit, Network, and Security (SANS) Institute - Glossary of Security Terms,

Steve Launius

- Information Systems Audit and Control Association (ISACA) - Cybersecurity Fundamentals Glossary,
- International Organization for Standardization (ISO) - Search for Terms & Definitions,
- Internet Engineering Task Force (IETF) Trust - Request for Comments (RFC) 4949 Internet Security Glossary,
- Information Technology Infrastructure Library (ITIL) v3 - Foundation Course Glossary.

There is some agreement between definitions, but it is not reasonable for non-technical professionals to learn the abundant terms and nuances of each. A smaller set of organizational-wide IT threat terms are necessary for more business-oriented professionals.

A discrete set of IT threat categories with standard definitions can increase communication and support risk reduction. Information security operations provide analysts with a rich vocabulary of cyber threat terms and a structure for appropriately characterizing attacks. CTI and incident response operations describe and analyze an attack in great detail to support threat hunting, sharing, and governance of information security operations. A taxonomy of IT threat terms can provide appropriate categories at various levels of granularity to aid threat analysis, risk assessments, and ultimately decision-making. Capturing and organizing unstructured threat information through CTI and incident response activities requires a standard set of threat terminology. Reports and metrics with a common set of terms can speed comprehension of the threats and incident response times. Business unit management and organizational leadership can more quickly understand the greatest threats to their organization after reviewing threat reports and metrics with standard terminology.

Since organizational leadership makes decisions based on risk, threat terms must be able to support risk management. All businesses must balance risk with reward, but severe consequences may result from misunderstanding the risk. An accurate depiction of the threats to information technology is vital for leadership to make appropriate decisions. Organizations in many industries use a variety of risk frameworks that may be threat-, vulnerability-, or asset-based. Regardless of the risk framework type, the

Steve Launius

quantities of threats should be commensurate with the maturity of the organization's risk management. Listing every possible hazard in an immature implementation of a risk framework can overwhelm risk analysis and bring the process to a halt. The risk management process should use threat categories appropriate for the maturity of the organization's risk assessment.

2.2. Threat Taxonomy for Cyber Threat Intelligence

CTI was born from the application of military intelligence doctrine to data analysis of cyberattacks. The DOD describes the intelligence process as a cycle of phases: direction, collection, processing, analysis, dissemination, and feedback (JP 2-0, 2013). While represented as a cycle, the steps may happen concurrently or may be skipped entirely depending on the situation. The intelligence cycle prescribes the process for collecting threat data and transforming it into threat intelligence. Brian P. Kime's article, "Intelligence Preparation of the Cyber Operational Environment" relates the DOD intelligence cycle to information security by presenting a collection method for threat data from IT infrastructure (Kime, 2016). Figure 1 shows the transformation of threat data into information, via structure and context, then into intelligence, via analysis, as it flows through the intelligence cycle phases.

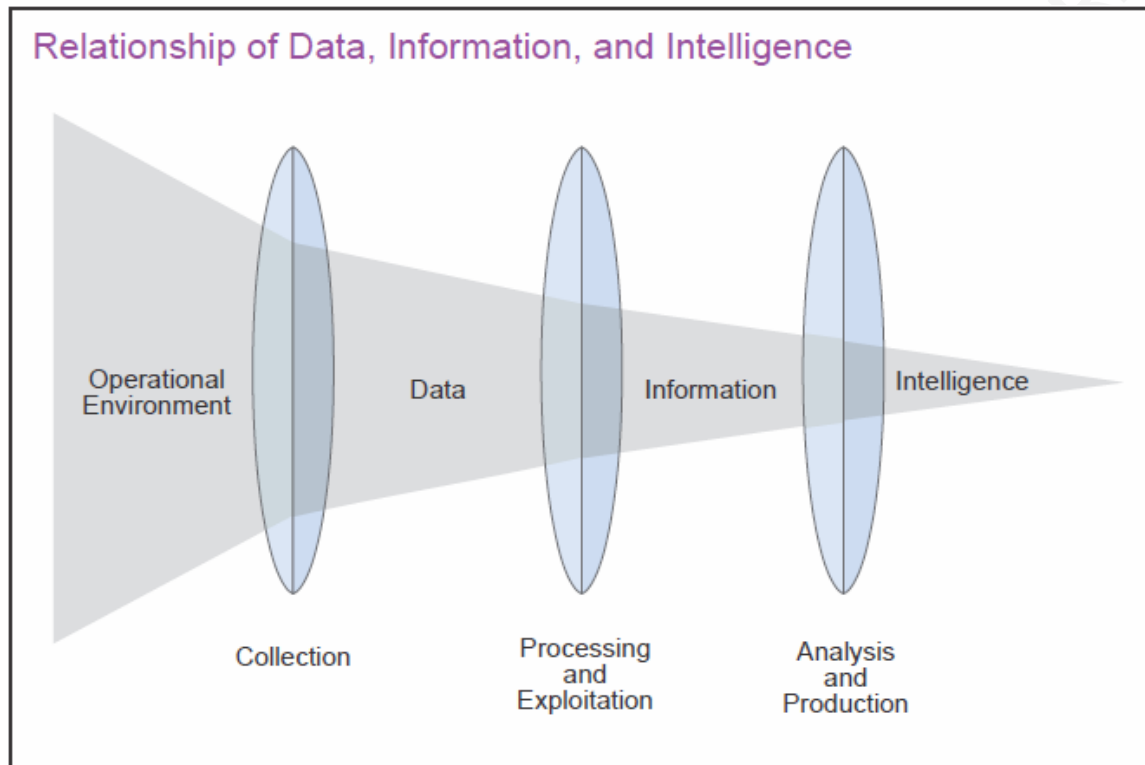


Figure 1 Relationship of threat data, information, and intelligence.

Structuring data to produce information is precisely where an IT threat taxonomy fits into CTI. A threat taxonomy sits on top of the available standards and ontologies for capturing threat data.

There are several CTI standards for modeling, storing and sharing threat data from cyberattack investigations. These standards capture indicators of compromise (IOC) or attacker tactics, techniques, and procedures (TTP). IOC are the easy to modify artifacts with the context pertinent to a cyberattack, such as file hashes of malicious program files or domain names of phishing websites. TTP describe the actions, skills, methods, or modus operandi (MO) an adversary uses to accomplish their goals. Threat models help relate IOC and TTP to each other for an illustration of the overall attack process and objectives during analysis. Robert M. Lee and Mike Cloppert describe threat modeling, such as Cyber Kill Chain and Diamond models, as an intrusion analysis technique for understanding threats and prioritizing defensive efforts that drive security (Lee, 2016). Organization and collection of the similar actions and techniques of cyberattacks facilitate sharing between industry partners and government bodies. Greg Farnham's paper on "Tools and Standards for Cyber Threat Intelligence Projects" (Farnham, 2013)

Steve Launius

presents and defines many CTI standards for an evaluation of a project management process. Those relevant for storing and sharing TTP include Structured Threat Information eXpression (STIX), Open Indicators of Compromise (OpenIOC) framework, and Collective Intelligence Framework (CIF).

While CTI standards provide structure for comprehensive threat analysis by subject matter experts, they often lack general groupings necessary for decision-makers to understand threats. According to the SANS 2017 CTI Survey (Shackleford, 2017), CTI standards have seen widespread adoption within CTI programs since Farnham's article was published. STIX consists of even more granular CTI standards. The Common Attack Pattern Enumeration and Classification (CAPEC) is a standard for describing cyberattack patterns (MITRE, 2017) that fits into STIX. CAPEC has 508 terms to portray all possible attack patterns. STIX and CAPEC are examples of the intricate threat detail capable with CTI standards. These capabilities aid threat analysis, but a higher-level perspective supports strategic CTI products.

CTI has three levels of analysis with a different purpose and audience for each: strategic, operational, and tactical. The operational and tactical levels of intelligence analysis concentrate on tracking and sharing attacker IOC and TTP with the CTI standards as previously explained. Analysis at the strategic level of CTI requires the same threat information, but addresses the overall risk to the organization by answering questions about cyber threats from leadership. The "Operational Level of Cyber Intelligence" published in the *International Journal of Intelligence and CounterIntelligence* provides an overview of these levels suitable for this discussion (Mattern, 2014). Strategic level intelligence "... pertains to an organization's general direction, specific goals, and resource allocation in service of its mission, as guided by the highest-level executive or command entity." Strategic intelligence analysis includes comparing security resources to trend changes in threats over time. At this level, intelligence analysis informs business units about the most likely threats to impact operations and the resources necessary to reduce this risk. A threat taxonomy supports strategic intelligence analysis with a consistent threat perspective to satisfy the needs of organizational leadership.

Steve Launius

Within the private sector, CTI operations concentrate on operational and tactical levels of analysis. The SANS Institute sponsors an annual survey of CTI since 2015 that demonstrates a focus on operational and tactical intelligence analysis, specifically on IOC. Comparison of the last three reports reveals a growing adoption of CTI with security tools primarily designed for identification, collection, or correlation of IOC. According to the 2015 survey, CTI improves security and response by increasing visibility into attack methodologies, cited by 63% of respondents, and by increasing incident response times, cited by 51% of respondents (Shackleford, 2015). The top three use cases in the 2016 survey were blocking malicious IP addresses or domain names at the firewall, adding context to incidents, and identifying malicious activity through DNS logs (Shackleford, 2016). The 2017 survey indicates that most organizations have dedicated CTI teams for collecting and processing CTI data (Shackleford, 2017).

These same studies also show the lack of application to strategic analysis. In the 2016 survey, more than half of the respondents said CTI is important to risk prioritization and decision making, but the 2017 survey lists “budget and spending prioritization and decisions” lowest among the use cases for CTI. Therefore, it is not surprising to see the primary skills for strategic analysis reporting, writing, presentation, and oral communications at the bottom of the skills list for CTI analysts in the 2017 survey. The survey respondents indicate the value of CTI is from an increase in preventing attacks and responding to attacks. However, CTI does not appear to be affecting strategic-level decisions. An inability to communicate with business terms the sources threatening specific business operations and the appropriate security measures to reduce this risk are the likely reasons why CTI is not influencing leadership.

Standard threat categories and terms in a taxonomy of all IT threats can assist analysis for producing strategic-level intelligence. Many publicly available intelligence sources produce unstructured reports. These intelligence sources frequently describe the same threat with various synonyms or attack terms. There is little agreement between sources of the names given to adversaries, malware, or attack techniques. Aggregation of the threat components, while consuming intelligence from a variety of sources, supports automated analysis methods. A threat taxonomy can help match these external reports to internal incidents. Organizations can predict future adversary actions by identifying

Steve Launius

attack patterns when threat modeling has a standard terminology. Revealing trends in attack vectors and adversary methods is possible when analyzing cyberattacks with a threat taxonomy. This type of analysis is useful for risk management because identifying the most likely threats helps prioritize remediation. Threat frameworks with detailed ontologies of threat information are difficult to use in risk analysis. Given the number of possible actors, actions, targets, and consequences for every threat, the list of possible threat events may total in the thousands or more. Governance, risk, and compliance (GRC) tools can provide an organization with automation of risk assessment calculations for complex threats. However, GRC tools are not available within every organization or may not support CTI standards. In the absence of these tools, scripts or macro-enabled productivity software can provide sufficient automation of workflow to produce CTI products usable in a risk assessment. Grouping threat information into a taxonomy provides a finite set of threat scenarios, so the risk analysis process does not overwhelm available resources.

2.3. Threat Taxonomy for Risk Assessments

The rich threat information in CTI can support information security risk frameworks, but assessing non-adversarial threats is also important. An adversarial threat taxonomy in a CTI program needs to be merged with non-adversarial threats, like environmental or human mistakes, in a risk assessment to communicate the level of risk across all threats facing an organization's information services. Risk frameworks from organizations like NIST, ISO, US-CERT, ISACA, and others use likelihood estimates for both adversarial and non-adversarial threats in the assessment process. The Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) is a risk management methodology from Carnegie Mellon University and US-CERT. OCTAVE Allegro (the most recent version) is an information asset-based assessment methodology which uses simple qualitative assessments of threat profiles. ISACA's latest version of Control Objectives for Information and Related Technologies (COBIT) 5 for Risk addresses risk of enterprise IT governance in the form of principals and guidance. To demonstrate the integration of a threat taxonomy into a risk framework the NIST's Risk Management Framework (RMF) provides a useful open and mature framework (NIST SP 800-37,

Steve Launius

2010). NIST's Guide for Conducting Risk Assessment (NIST SP 800-30, 2012) provides important concepts and processes for implementing the RMF and describes where a threat taxonomy interacts with the risk assessment. Identifying, estimating, and prioritizing information security risks are the function of a risk assessment.

Threats are one common risk factor NIST's risk assessment methodology identifies for assessing and relating risks in a model. The risk factors define the characteristics for determining risk levels that are essential for communicating problematic situations. Definitions for risk factors are informed by an organization's risk management strategy or during risk framing if a strategy does not exist. The other key risk factors seen below in Figure 2 include vulnerability, impact, likelihood, and predisposing condition. Threats break down into *threat sources* that cause *threat events*. A threat event has potential to negatively impact an organization's operations or assets through the loss of confidentiality, integrity, or availability of information or information systems. A threat source is the "intent and method targeted at the intentional exploitation of a vulnerability or a situation and method that may accidentally exploit a vulnerability" (NIST SP 800-30, 2012). NIST's comprehensive overview of threat sources includes:

- Cyber or physical attacks
- Human errors
- Failure of resources
- Environmental disasters, accidents, or failures

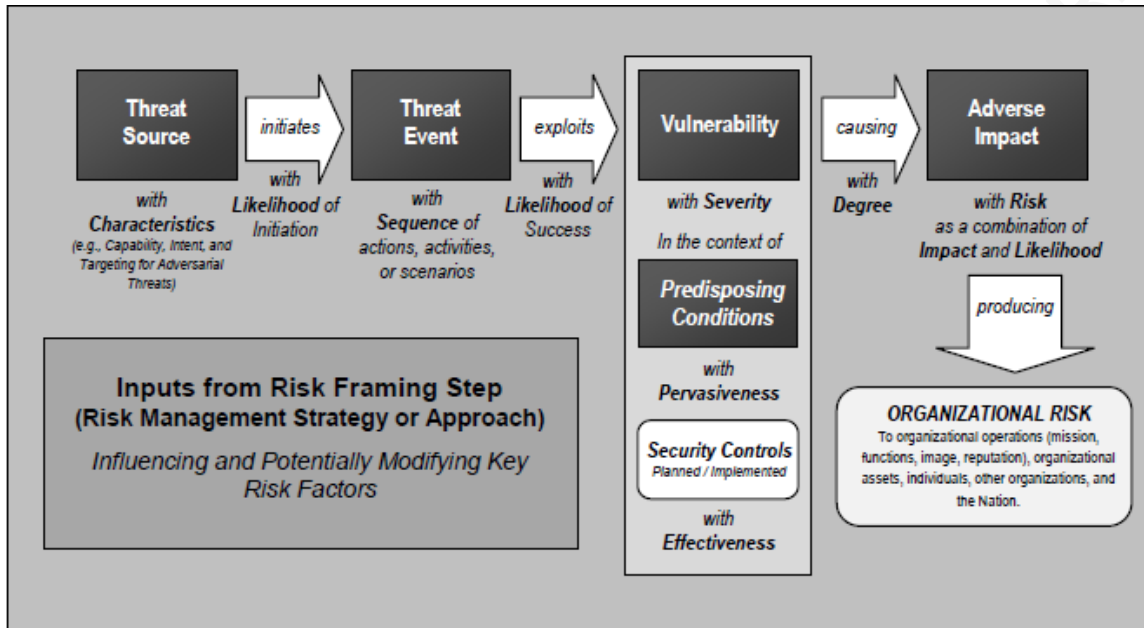


Figure 2 NIST 800-30 generic risk model with key risk factors.

NIST prescribes a four-step risk assessment process, illustrated in Figure 3, for preparing, conducting, communicating results, and maintaining a risk assessment. Organizations define and use the threat taxonomy in the first two steps of the risk assessment process. During Communicate Results in the third step, the report and metric products sent to leadership should use this same threat terminology.




Figure 3 Risk assessment steps from NIST 800-30.

In preparation for the risk assessment, organizations can define a threat taxonomy in the first step as part of risk framing. Identifying the main assumptions relevant to risk assessments is one of the tasks which enables the RMF to clarify risk models and increase repeatability of results. Two of the key assumption areas are threat sources and events. The level of detail chosen for threat sources and events will establish the set of possible threats available when identifying the relevant threats to the organization in the Conduct Assessment step.

Another crucial assumption area for risk assessments is the analytic approach for characterizing threat sources and events. The analytic approach consists of both the assessment type (i.e. quantitative, qualitative) and analysis type (i.e. threat-, asset-, of vulnerability-orientated). A many-to-many relationship exists among threat events and sources, therefore levels with greater detail increases the complexity of the risk assessment. A threat taxonomy categorizing all possible threat sources and events with varying levels of granularity can allow an organization to move from less to more detail as their risk management program matures.

Steve Launius

3. Comprehensive Threat Taxonomies

A taxonomy is an ordered classification system, often hierarchical, where each parent tier is a grouping of the terms characterizing its child tier. The terms each taxonomy uses for the hierarchical levels are slightly different but serve a similar purpose. Descriptive terms for the top-level of a taxonomy may include class, top-tier, or high-level. Terms for the second level of a taxonomy may include family, threats, or subclasses. The designations for taxonomies with a third level consist of elements or threat details. The terms and structure of each taxonomy used in this research can be found in Appendix A.

Several institutions have created comprehensive threat taxonomies for IT systems. A comprehensive threat taxonomy will have several features. A simple hierarchal structure is necessary where the top-tier has no more than ten categories. This discrete set of categories must work to organize events, activities, situations, or contexts from diverse sources of threats encompassing both adversarial and non-adversarial threats. The taxonomy will only categorize the threat event component, but events must include activities from both human and environmental threat sources. The subcategories should include more detail than the higher-level groups with definitions for the terms. Definitions of all threat categories are valuable for creating consensus among the professionals who will work with the taxonomy.

Most of the qualifying taxonomies are incomplete as work on them has only begun within the last few years. Each taxonomy has a different goal and purpose that shapes the categories selected for it. For example, the business operational threat categories of Carnegie Mellon University's taxonomy use business-orientated terms including *people*, *process*, *technology*, and *external*. Mapping these taxonomies should be straightforward with any of the published security control recommendations, like NIST 800-171. The threat taxonomies are primarily for organizations with threat intelligence capabilities to provide probability estimates for threat activities during risk management. In addition to a review of the goal and purpose of the taxonomies, a short analysis of their qualities will reveal their strengths and weaknesses.

Steve Launius

3.1. Open Threat Taxonomy

The goal of the Open Threat Taxonomy (OTT) was to create a shared and comprehensive set of information system threats that organizations may face. James and Kelli Tarala, authors of the OTT and owners of the security firm Enclave Security, released version 1.1 as an open source tool in October 2015. The OTT defines a threat as “... the potential for a threat agent to cause loss or damage to an information system” (Tarala, 2015). Part of the complexity of defining threats comes from the components that compromise a threat. The OTT lists these components as threat source or agent, threat action, threat target, and threat consequence. Tarala describes the relationship of these components as, “A threat source will most often perform a threat action against a threat target, which leads to threat consequences” (Tarala, 2015). This taxonomy only describes threat actions, but uniquely includes a priority ranking for each action. A one to five scale ranks the priority of each threat, where priority should go to threats with a higher rank. Threat models and attack observations from contributors to the OTT help establish the priority scores and “should be viewed as consensus guidance” (Tarala, 2015).

The OTT covers most of the pertinent threats to information system operations without forgetting most of the non-technical dangers. The OTT categorizes threats by their nature and by the extent to which they impact the confidentiality, integrity or availability of information systems. This taxonomy has a total of 75 threat actions broken down into four main categories:

- Physical Threats
- Personnel Threats
- Resource Threats
- Technical Threats

Definitions for each category elaborate on the nature of each threat group. However, the threat actions do not have definitions, only clear descriptive terms. Even though there are short action phrases, an audience’s experience could lead to ambiguous interpretations of the terms. The small set of threat categories describes actions that can cause damage to information systems. Adverse impact is defined as threats to confidentiality, integrity, or availability of each category. Therefore, many of the threat actions have an adversarial perspective. This grouping perspective results in a concentration of threat actions within the Technical Threats category as technical vulnerabilities in information systems are

Steve Launius

numerous. The categorization of all possible threat sources is incomplete, as capturing legal threats does not appear to be possible in the OTT.

The holistic coverage of information systems threats from OTT can provide broad risk comparison across an organization. The OTT works well with risk frameworks that consider inherit and residual risks separately. This is due to priority ranking scores a group of industry experts assigns to each OTT threat action. This ranking system allows an organization to prioritize one threat over another when it must choose between investing in resources to mitigate threats with the same likelihood of occurring. Besides the threat actions, the taxonomy does not address other threat components or help with identifying mitigation controls. Mapping the threat actions to specific security controls, such as NIST 800-53, could assist in completing a risk assessment.

3.2. ENISA Threat Taxonomy

In January 2016, the European Union Agency for Network and Information Security (ENISA) published a taxonomy as an aid for threat information collection and consolidation (ENISA, 2016). The ENISA Threat Taxonomy (ETT) defines Cyber Threats as "... threats applying to assets related to information and communication technology." ENISA's purpose for its taxonomy is to provide definitions for threat terms with a possibility of rearranging its structure. The ETT was designed as an analysis mechanism for collecting and sorting threat information.

The ETT provides a unique view of possible threat actions, but without the consistency and clarity found in other taxonomies. The eight or nine, depending on the version, high-level categories of the ETT are a mixture of consequences and intentions for the 75 total threats actions. The *high-level threats* include:

- Physical Attack
- Unintentional Damages
- Disasters
- Failures / Malfunction
- Outages
- Eavesdropping / Interception / Hijacking
- Nefarious Activity / Abuse
- Legal

The *threats* and *threat details* make up the next two levels of the ETT creating one of the most detailed threat taxonomies. While there is an expectation of change for different

Steve Launius

versions of a taxonomy, the lack of consistent relationships and accurate definitions throughout the ETT detract from the purpose of a taxonomy. One inconsistency is the alternate terms for three of the high-level threats. The ETT uses a slash symbol to expand the terms of these categories instead of using a single term and definition like the other categories. The high-level threat definitions do not support mutually exclusive categories. For example, the Eavesdropping threat has a definition that fits into the Nefarious Activity threat, but these categories exist at the same level. Additionally, several of the threats and threat details include the threat source or intentions in the description restricting its scope, which will lead to necessary revisions in the future. The lack of delineation between threat events and sources also causes ambiguous classification of a threat into multiple categories. Such a classification supports complex relationships in threat ontologies, but conflicts with the simplifying purpose of a taxonomy. Similar to OTT, the ETT adversarial threats focus on attacker actions that can negatively impact information systems but disperses them into more high-level threat categories. The ETT brings legal threats clearly into consideration with the inclusion of a Legal category for regulations, changes in law, and the political environment.

3.3. NIST Risk Assessment Threat Exemplary

The appendix within NIST's Guide for Conducting a Risk Assessment includes exemplary threat events that provide a sample threat taxonomy. NIST's risk model decomposes threats into a source and event for analysis of a single threat. A series of threat events can create a threat scenario that NIST defines as "a set of discrete threat events, attributed to a specific threat source or multiple threat sources, ordered in time, that result in adverse effects" (NIST SP 800-30, 2012). Multiple events from the same threat source or multiple threat sources executing the same threat event may compromise threat scenarios. These scenarios can result in many granular circumstances; therefore, a mature risk management process is necessary to handle the numerous scenarios that result from this analysis. An organization need only to assess the relevant threat events when there is an adversary with intent or capability to initiate an attack.

Steve Launius

For consistent comparisons with other taxonomies, the evaluation will only include the NIST exemplary threat events. The NIST model breaks all threat events into two high-level categories:

- Adversarial
- Non-adversarial

The two-level hierarchy in this taxonomy results in a concentration of threat events for the adversarial category. The second-level categorizations of adversarial threat events are similar to the stages in the Lockheed Martin kill chain model (Lockheed) that characterize adversarial TTP. These stages of a cyberattack include reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on objectives. The NIST guide references the MITRE Corporation's CAPEC for characterizing cyberattacks with greater detail (CAPEC, 2017). These adversarial attack patterns describe possible methods for exploiting information systems from an attacker's perspective. The adversarial events categorized by the kill chain stages can be useful for mapping with security controls, like NIST SP 800-53. There are far fewer non-adversarial threat events in NIST's taxonomy and, therefore, no additional subcategories for this type of threat. The non-adversarial category is also lacking many of actions found in other taxonomies for unintentional, accidental, legal, or other non-malicious actions. This sample threat taxonomy may not be useful for an organization unless the threat categories are extended.

3.4. Taxonomy of Operational Cyber Security Risks

A comprehensive threat taxonomy from Carnegie Mellon University is one of the oldest available. In 2010, the Software Engineering Institute (SEI), a federally funded research and development center based at Carnegie Mellon, produced the first version of the Taxonomy of Operational Cyber Security Risks (TOCSR) (CMU/SEI, 2014). The taxonomy was updated in 2014 to map with the security and privacy controls in Version 4 of NIST SP 800-53. This taxonomy categorizes instances of *operational cyber security risks* defined as “operational risks to information and technology assets that have consequences affecting the confidentiality, availability, or integrity of information or information systems.” The purpose of TOCSR is to provide a tool for identifying all the operational cyber security risks within an organization.

Steve Launius

The concise terms and categorization method of TOCSR produces a taxonomy that can assist in risk assessment activities. The primary emphasis of the categorization method is on operational risks to information systems. The TOCSR characterizes threats from a business risk perspective, instead of a threat source perspective as in the other threat taxonomies. This results in categories of threats actions for people, process, and technology. This method results in four top-level categories that SEI calls classes:

- Actions of people
- Failed internal processes
- Systems and technology failures
- External events

In SEI's terminology, each class decomposes further into subclasses and elements. The operational risk terms from Risk Lexicon from DHS (DHS, 2008) are the basis for the threat categories. While this taxonomy aligns with SEI's OCTAVE method for risk assessments, threat taxonomies are not exclusive to one risk framework. Representation of a complete attack scenario may require a combination of TOCSR threat categories. For practical implementation in the NIST risk assessment, threat elements from multiple classes or subclasses will compose a single scenario. For example, a *software* flaw present in a production web application due to inadequate *testing* could be a result of any element under *actions of people*. SEI provides a mapping to the security guidelines in NIST 800-53.

3.5. Other Threat Taxonomies

There are several other published taxonomies for adversarial threats or intelligence sharing. As the need for a taxonomy arose with the formal gathering and sharing of cyberattack information, the work of developing suitable taxonomies is still ongoing. Many organizations only address the most prevalent threats or create taxonomies for specific threats. In either case, these taxonomies are not suitable for an organization-wide taxonomy of threats.

There are many more adversarial-centric threat taxonomies which provide a multitude of options for categorizing the variety of malicious human cyber activities. However, these do not allow comparisons with environmental threats and therefore do not meet the criteria for consideration of a comprehensive threat taxonomy. The aforementioned CAPEC is one such taxonomy of cyberattack patterns by MITRE.

Steve Launius

Another adversary-centric taxonomy comes from the US government called the Cyber Threat Framework (CTF). The CTF was designed to improve communication between cyber experts and senior leadership across many departments throughout the intelligence community (ODNI, 2017). The variety of threat models in use at different government agencies made sharing cyber threats difficult because of different terminology that was highly technical. Many other CTI standards can map into the four stages of adversary cyberattacks in the CTF. The Office of Director of National Intelligence provides a lexicon for the CTF that equates to a threat taxonomy. The flexible design of the framework allows different views of same adversarial threat information for diverse audiences. One final example of an adversarial threat taxonomy comes from Agari, a secure email exchange company, specifically for cyberattacks against messaging systems (Jakobsson, 2017). The taxonomy breaks down the steps for attacking an email system that was extended to all types of messaging systems, including instant messaging. The scope of these adversarial threat taxonomies is too narrow for organizing a comprehensive set of threats meant for an organization-wide risk assessment.

Researchers at Georgetown University are creating a taxonomy for the existing threat intelligence sharing standards. This cyber threat intelligence information sharing exchange ecosystem (CyberISE) (Burger, 2014) is a classification system for CTI sharing standards. Eric Burger's research presents the structure and relationship to other information sharing technology. The organization of the CyberISE has five top-level categories in a layered model, mimicking the Open Systems Interconnection (OSI) model. The two lower layers address the exchange and authorization of information sharing, while the three upper layers categorize the information exchange. The *Indicators* layer holds the details of an incident or cyberattack. The *Intelligence* layer contains actions to perform when detecting indicators or assessing threats. The *5W's* layer comprises the types of questions to ask incident indicators to determine whether an attack is occurring. Since the CyberISE model is for characterizing the existing information sharing standards, it is not an appropriate taxonomy for the categorization of threat information.

The Cambridge Risk Framework is a global threat taxonomy for business operations by the University of Cambridge. The report *A Taxonomy of Threats for*

Steve Launius

Complex Risk Management (Coburn, 2014) presents the Cambridge Taxonomy as a taxonomy of macro-catastrophe threats. The basis for threat categorization is extreme events with potential to cause damage or disrupt global social and economic systems. Extreme events have a large impact on global trade and commerce across multiple continents.

Cambridge's development methodology includes a review of historical events and disaster catalogs to create a hierarchy structure of 5 primary classes, 11 families, and 55 types. The report includes definitions for the five classes: Finance & Trade, Geopolitics & Society, Natural Catastrophe & Climate, Technology & Space, and Health & Humanity along with their corresponding families. Insurance risk management is a primary application of the Cambridge Taxonomy. Secondary functions involve risk management of business operations, national security, and finances. While extreme events will have some impact even to small business operations, the likelihood of a global macro-catastrophe event occurring should be overshadowed by more likely, local catastrophes for most businesses. Additionally, the other selected comprehensive threat taxonomies are IT-centric to the effects of threat events. Therefore, the Cambridge Taxonomy was not included in this research evaluation, but global organizations may want to consider it. Organizations of any size may choose to consider this threat taxonomy by redefining catastrophes and extreme events to include disasters at any scale.

4. Threat Taxonomy Evaluation

This research evaluation of threat taxonomies uses a qualitative research survey. A qualitative research methodology best supports results dependent upon personal opinions and diverse perspectives. The primary survey focuses on a large financial services company. The risk management department of this company agreed to receive the survey. Responses from this source were plentiful with a total of 61 respondents, labeled as 'Financial Company' in the analysis. An attempt was made to obtain diverse perspectives outside of the Financial Services industry by posting the survey to several social networking forums including information security and educational email list serves as well as professional networking websites. Unfortunately, the response from these sources was much smaller with a total of 23 respondents, labeled as 'Non-Financial

Steve Launius

Company' in the analysis. The survey began by asking all respondents their industry and job role. To represent different perspectives the analysis compares responses from four groups: Management, Non-Management, Financial Company, and Non-Financial Company. Presentation of the terms and structure of each taxonomy were straightforward, but minor changes were necessary due to formatting restrictions in the survey tool.

There is a potential for respondents to favor the presentation format of a taxonomy while presenting the survey. Authors of the taxonomies use various formatting styles in publications, but to avoid any bias the survey has a consistent table formatting for all the taxonomies. Presentation of the taxonomies took the form of uniform tables. The top-tier categories are set in header rows with the same blue color background. The second tier follows in the next row with categories in a bold font and specific threat actions in a bulleted list for the third tier. The survey mitigates further bias by presenting the taxonomies in a randomly chosen order.

The survey includes only the first two levels of the more complex taxonomies to keep respondent review time to a minimum. Both NIST and ENISA have three or more tiers that can be both overwhelming and tedious to review. The top two tiers list all the major threat categories for each taxonomy. However, the taxonomies presented without the third tier are likely to have lower ratings for completeness. This effect can be even more profound when the clarity of the top tier categories is low, indicating a respondent would not be able to infer the types of threats in a category without them explicitly listed. Reducing the threat actions in the OTT was also necessary for repetitive actions using similar methods. For example, reducing the eleven Application Exploitation actions with different attack methods into a single threat action in the Technical Threat category saves review time without detracting from the threat event. The length of the taxonomies was a likely factor in completing the survey. Fifteen percent of the respondents failed to complete review of all four taxonomies. The OTT had the most responses with about ten more than the other taxonomies. See Appendix B for a complete view of each taxonomy in the same presentation format and order.

The characteristics chosen for evaluation include completeness, complexity, and clarity. These traits were chosen for evaluation because they are ubiquitous, descriptive

Steve Launius

words and encompass the individual characteristics that make a taxonomy a useful tool for communication. Therefore, respondents did not receive definitions for the traits. The rating score for each of these characteristics consists of a weighted scale from 1 to 5, from worst to best, with the following common descriptions: Not at all, Slightly, Moderately, Quite, Extreme. The weighted answers provide a quick method for scoring and comparing the taxonomies.

A consistent analysis method compares results for each of the traits without favoring one over another. However, organizations may choose to favor one trait over another because of its available resources. An organization may find the clarity of threat terms more advantageous than completeness, for example, if there is no intranet website for sharing a central glossary and training employees is unlikely. On the other hand, favoring clarity may also imply favoring the least complex taxonomy, and vice versa, given the relationship between these two traits.

4.1. Completeness

A complete threat taxonomy would be able to characterize all possible threat actions or events. The categories chosen by a taxonomy may preclude certain types of threats. For example, the NIST non-adversarial categories do not incorporate threats from legal action. For each taxonomy, respondents were asked to select one rating for the completeness of the taxonomy from these answers (with weight): Not at all complete (1), Slightly complete (2), Moderately complete (3), Quite complete (4), or Extremely complete (5). The score calculation is the average sum of weighted responses for each group. Therefore, groups with higher values in Figure 4 indicate more responses and the completeness scores in each group rank each taxonomy.




Figure 4 Completeness scores of each threat taxonomy by respondent groups.

The overall completeness scores indicate OTT is the most complete. However, the Management group scores ENISA as the most complete. ENISA's taxonomy has the most threat actions present in the survey. Therefore, respondents may have given higher scores to ENISA based on this overall number. This is the most likely conclusion because respondents expect surveys to be brief. The low scores given to NIST further support this conclusion. NIST has the lowest number of threat actions in the survey because the length of adversarial threat actions in NIST SP 800-30 prevented listing them all in the survey application. The threat descriptions in NIST's adversarial tier create a cumbersome taxonomy table that is many pages long. The taxonomy review in previous sections shows that both NIST and OTT were unable to categorize legal threats. Additionally, NIST lacks more nuance for non-adversarial threats found in the other taxonomies. Even though scores for the TOCSR rank it third overall for completeness, the review in an earlier section did not find any events unfit for its threat categories. The business-risk perspective was likely a factor in lower completeness scores given its unique viewpoint from actions or failures of people, process, technology, or externalities.

4.2. Complexity

A complex threat taxonomy is one that is difficult to understand without additional context. Complexity could refer to either the structure or terms. Respondents may consider a taxonomy more complex if it has many high-level categories or more

Steve Launius

threat terms describing an event. For each taxonomy, respondents were asked to rate the overall complexity of each from these answers (with weight): Not at all complex (5), Slightly complex (4), Moderately complex (3), Quite complex (2), or Extremely complex (1). Score calculations follow the same process as in the completeness section. However, reversal of the weights is necessary to designate less complexity as the more desirable trait. Therefore, taxonomies with higher scores in Figure 5 are less complex.




Figure 5 Complexity scores of each threat taxonomy by respondent groups.

Respondents score the OTT and TOCSR as the least complex taxonomies. These taxonomies both have four top-tier categories with the most concise terminology to describe threat actions. The Financial Company and Management groups score TOCSR complexity just below the OTT. These groups are more likely to have a business-centric perspective that contributes to rating TOCSR higher than the other groups. However, these groups still rate the OTT as the least complex. Along with the Non-Management and Non-Financial groups both rating the OTT as the least complex, by larger margins, the overall score makes it the least complex taxonomy.

4.3. Clarity

A clear taxonomy would have simple threat terms and threat events that are logically relevant under the same category. While definitions are an essential element of a taxonomy for maintaining consistency, simple threat terms should plainly characterize a common set of threat events. For each taxonomy, respondents were asked to select a

rating for the clarity of terms from these answers (with weight): Not at all clear (1), Slightly clear (2), Moderately clear (3), Quite clear (4), or Extremely clear (5). The score calculation is the average sum of weighted responses for each group. Thus, the clearest taxonomies in Figure 6 have a higher score.




Figure 6 Clarity scores of each threat taxonomy by respondent groups.

All the respondent groups rate the OTT as the clearest taxonomy. Only in the Management group did both the ENISA and TOCSR taxonomies have clarity scores close to the OTT. The respondent groups rate ENISA second, or a close third, in clarity. High clarity scores for ENISA's taxonomy were unexpected because of its alternative terms for several categories. Although, respondents may have seen the alternative terms as more descriptive characteristics for a category. Even though the TOCSR has the most concise terms for threat actions, its business-risk perspective appears to have detracted from the overall understanding of the terms by respondents.

4.4. Overall

The Open Threat Taxonomy overall scores are the highest for the completeness, complexity, and clarity traits. The combined group scores for each trait are viewable side-by-side in Figure 7. While the overall preference is for the OTT, both ENISA and TOCSR have strengths in different traits. The TOCSR has a high score for complexity, and the completeness score for ENISA is high. An organization favoring complexity or

completeness may also consider either of these taxonomies. However, when it comes to clarity, the OTT outscores the other taxonomies by a large margin of at least ten points.

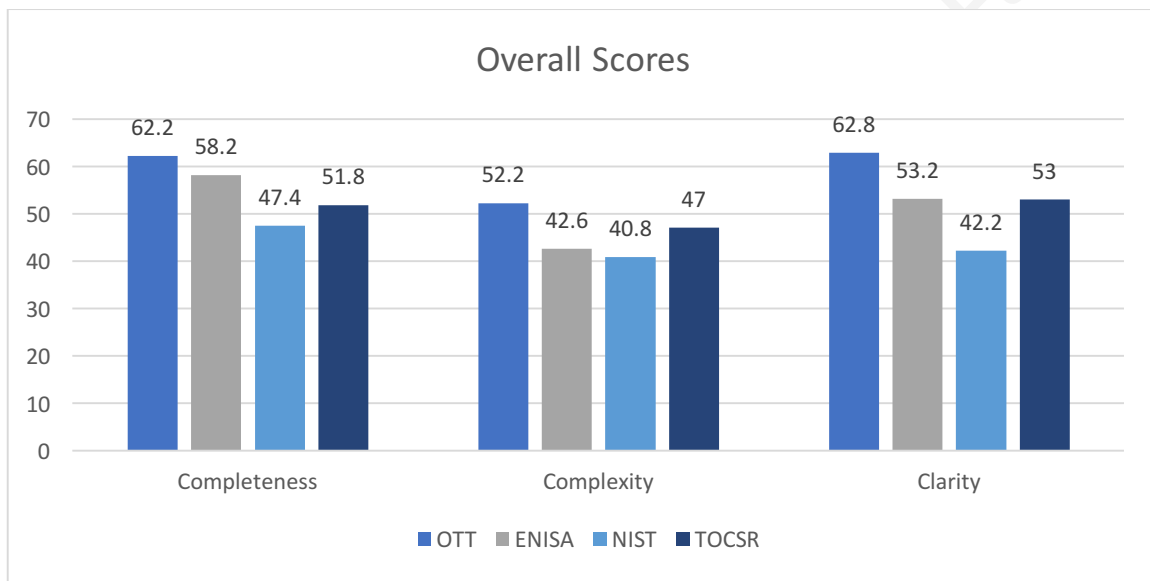


Figure 7 Overall scores of each taxonomy by traits.

5. Conclusion

Survey respondents were asked to rate the clarity of terms to determine which threat taxonomy had the simplest terms and most logical grouping. Simple terms can help an organization's leadership understand threats to operations dependent on information technology. Many threat terms are available in CTI standards for intrusion analysis. However, there are too many terms for non-technical decision-makers to understand. Additionally, threat categories that logically group similar terms are clearer.

Review of the structure and terms of each threat taxonomy in the survey allowed respondents to judge which is the least complex. The exhaustive detail and multiple relationships within CTI standards that make them good for intrusion analysis also make them a poor choice for communicating with leadership. A smaller set of threat categories can reduce the complexity of cyberattacks for this audience. Grouping threat events with a hierarchical system can also reduce complexity when each category has similar events. The multiple levels within a hierarchical taxonomy provide several granularity options. This structure allows an organization to use the appropriate level for its risk assessment as it matures. Higher levels can help keep the risk assessment simple and small when it is

Steve Launius

immature. Lower levels can provide greater detail for complex threat scenarios when the organization is ready.

In order to assess the degree of inclusiveness for each threat taxonomy, the survey inquired about the completeness. Cyberattacks are not the only threats to an organization's information technology. Threats may arise from natural disasters, legal discussions or political interests, or employee accidents. The CTI standards concentrate on an adversary's malicious activity, so the lexicon in these standards is missing terms that characterize alternative threat sources. Risk frameworks help model all types of threats facing an organization. Comprehensive threat taxonomies fit into risk assessments, like NIST SP 800-30, to present decision-makers with a risk comparison across all of the threats.

This research found several methods for categorizing all of the possible threats to information technology. Only a handful of these threat taxonomies attempted to address all potential threats to IT within an organization. These nascent threat taxonomies may not be inclusive of all possible threats. The most mature taxonomy is about eight years old and updates have been infrequent. Since threat actions are one of the primary inputs for assessing IT risk, a public consensus of all the threats to information technology can improve communication within and between organizations.

The evaluation by both management and non-management personnel of these threat taxonomies strengthens the results of this research. The opinions of these two groups are vital for different reasons. Management needs to understand threats to improve communications with analysts and other business units in order to make quick decisions that influence the security resources of an organization. Non-management needs to present the threats to management, so they might obtain the necessary resources to address increasing threats. A familiar set of threat terms in meetings, reports, metrics, and risk assessments can help improve this communication. Based on the rating given for completeness, complexity, and clarity, this evaluation suggests each group prefers the Open Threat Taxonomy. This threat taxonomy can provide a complete picture of threat actions, with clear terms, in a manner that is simple for an organization's leadership to understand.

Steve Launius

5.1. Future Research

This analysis resulted in the selection of a preferred threat taxonomy. However, this evaluation excludes an assessment of taxonomies to aid in decision-making by leadership. Evaluation of decision-making would require implementing a taxonomy into a risk assessment, mapping to security controls, and reviewing the issues which may arise from this implementation. Many of the risk frameworks present qualitative methods for assessments, but a quantitative assessment may favor one taxonomy over another. A comparative case study utilizing different threat taxonomies for threat scenarios with different risk frameworks, or the same risk framework with different assessment techniques are two possible evaluation ideas. Keys to success for this implementation would include mapping to security controls, like NIST SP 800-53, or security requirements, like NIST SP 800-171, and calculating probabilities of occurrence and impact based on changes to the threat landscape.

References

- Definition of Virus. (2018, February 11). Retrieved from <https://www.medicinenet.com/script/main/art.asp?articlekey=5997>
- Glossary of Security Terms. (2018, February 11). Retrieved from <https://www.sans.org/security-resources/glossary-of-terms/>
- US Department of Defense, Joint Chiefs of Staff. (2013, October 22). Joint Intelligence (Joint Publication (JP) 2-0).
- Kime, B. P. (2016). Threat Intelligence: Planning and Direction. *SANS Institute*. Retrieved from <https://www.sans.org/reading-room/whitepapers/threats/threat-intelligence-planning-direction-36857>
- Lee, R. M., Cloppert, M. (2016). Forensics 578: Cyber Threat Intelligence. The SANS Institute. www.sans.org/course/cyber-threat-intelligence
- Farnham, G. (2013). Tools and Standards for Cyber Threat Intelligence Projects. SANS Institute. Retrieved from <https://www.sans.org/reading-room/whitepapers/warfare/tools-standards-cyber-threat-intelligence-projects-34375>
- Shackleford, D. (2017). Cyber Threat Intelligence Uses, Successes and Failures: The SANS 2017 CTI Survey. SANS Institute. Retrieved from www.sans.org/reading-room/whitepapers/analyst/cyber-threat-intelligence-uses-successes-failures-2017-cti-survey-37677
- MITRE Corporation. (2017, December 8). CAPEC - Common Attack Pattern Enumeration and Classification (CAPEC). Retrieved from <https://capec.mitre.org/>

Steve Launius

Mattern, T., Felker, J., Borum, R., Bamford, G. (2014). Operational Levels of Cyber Intelligence. International Journal of Intelligence and CounterIntelligence.

Published Aug 6, 2014. Volume 27, Issue 4, Pages 702-719. DOI:

10.1080/08850607.2014.924811

Shackleford, D. (2015). Who's Using Cyberthreat Intelligence and How? SANS Institute.

Retrieved from [https://www.sans.org/reading-](https://www.sans.org/reading-room/whitepapers/analyst/cyberthreat-intelligence-how-35767)

[room/whitepapers/analyst/cyberthreat-intelligence-how-35767](https://www.sans.org/reading-room/whitepapers/analyst/cyberthreat-intelligence-how-35767)

Shackleford, D. (2016). The SANS State of Cyber Threat Intelligence Survey: CTI

Important and Maturing. SANS Institute. Retrieved from

[https://www.sans.org/reading-room/whitepapers/bestprac/state-cyber-threat-](https://www.sans.org/reading-room/whitepapers/bestprac/state-cyber-threat-intelligence-survey-cti-important-maturing-37177)

[intelligence-survey-cti-important-maturing-37177](https://www.sans.org/reading-room/whitepapers/bestprac/state-cyber-threat-intelligence-survey-cti-important-maturing-37177)

National Institute of Standards and Technology. (2010). Guide for Applying the Risk

Management Framework to Federal Information Systems: a Security Life Cycle

Approach. Special Publication (NIST SP) 800-37 Rev 1.

<http://dx.doi.org/10.6028/NIST.SP.800-37r1>

National Institute of Standards and Technology. (2012). Guide for Conducting Risk

Assessments. Special Publication (NIST SP) 800-30 Rev 1.

<http://dx.doi.org/10.6028/NIST.SP.800-30r1>

Tarala, J., Tarala, K. (2015). Open Threat Taxonomy version 1.1. Enclave Security.

Retrieved from

http://www.auditscripts.com/resources/open_threat_taxonomy_v1.1a.pdf

Steve Launius

- European Union Agency for Network and Information Security. (2016). ENISA Threat Landscape. ENISA. <https://www.enisa.europa.eu/topics/threat-risk-management/threats-and-trends/enisa-threat-landscape/enisa-threat-landscape>
- Hutchins, E. M., Cloppert, M. J., Amin, R. M. (n.d.). Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. Lockheed Martin Corporation (Lockheed). Retrieved from <https://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf>
- Cebula, James., Popeck, Mary., & Young, Lisa. (2014). A Taxonomy of Operational Cyber Security Risks Version 2 (CMU/SEI-2014-TN-006). Retrieved from the Software Engineering Institute, Carnegie Mellon University website: <http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=91013>
- Department of Homeland Security Risk Steering Committee. (2008, September) DHS Risk Lexicon. Department of Homeland Security (DHS). http://www.dhs.gov/xlibrary/assets/dhs_risk_lexicon.pdf
- Office of the Director of National Intelligence (ODNI). (n.d.). Cyber Threat Framework Version 4.0 Lexicon of concepts and definitions. Retrieved July 15, 2017 from https://www.dni.gov/files/ODNI/documents/features/Cyber_Threat_Framework_Lexicon.pdf
- Jakobsson, M. (2017). The Threat Taxonomy: A Working Framework to Describe Cyber Attacks. Agari. Retrieved on October 28, 2017 from <https://www.agari.com/threat-taxonomy-framework-cyber-attacks/>

Steve Launius

Burger, E.W., Goodman, M. D., Kampanakis, P., Zhu K. A. (2014). Taxonomy Model for Cyber Threat Intelligence Information Exchange Technologies. Proceedings of the 2014 ACM Workshop on Information Sharing & Collaborative Security. Pages 51-60. DOI: 10.1145/2663876.2663883

Coburn, A.W.; Bowman, G.; Ruffle, S.J.; Foulser-Piggott, R.; Ralph, D.; Tuveson, M.; (2014). A Taxonomy of Threats for Complex Risk Management. Cambridge Risk Framework series. Centre for Risk Studies, University of Cambridge.

Appendix A: Threat Taxonomy Details

Note: The content as presented below was edited for presentation in the research survey; see references for complete taxonomies with definitions.

OTT Threat Actions & Ratings

Physical Threats

- Loss of Property
- Theft of Property
- Accidental Destruction of Property
- Natural Destruction of Property
- Intentional Destruction of Property
- Intentional Sabotage of Property
- Intentional Vandalism of Property
- Electrical System Failure
- HVAC Failure
- Structural Facility Failure
- Water Distribution System Failure
- Sanitation System Failure
- Natural Gas Distribution Failure
- Electronic Media Failure

Resource Threats

- Disruption of Water Resources
- Disruption of Fuel Resources
- Disruption of Materials Resources
- Disruption of Electrical Resources
- Disruption of Transportation Services
- Disruption of Communications Services
- Disruption of Emergency Services
- Disruption of Governmental Services
- Supplier Viability
- Supplier Supply Chain Failure
- Logistics Provider Failures
- Logistics Route Disruptions
- Technology Services Manipulation

Personnel Threats

- Personnel Labor / Skills Shortage
- Loss of Personnel Resources
- Social Engineering of Personnel Resources
- Disruption of Personnel Resources
- Negligent Personnel Resources
- Personnel Mistakes / Errors
- Personnel Inaction

Technical Threats

- Organizational Fingerprinting via Open Sources
- System Fingerprinting
- Credential Discovery
- Misuse of System Credentials
- Escalation of Privilege
- Abuse of System Privileges
- Memory Manipulation
- Cache Poisoning
- Physical Manipulation of Technical Device
- Manipulation of Trusted System
- Cryptanalysis
- Data Leakage / Theft
- Denial of Service
- Maintaining System Persistence
- Manipulation of Data in Transit / Use
- Capture of Data in Transit / Use
- Replay of Data in Transit / Use
- Misdelivery of Data
- Capture of Stored Data
- Manipulation of Stored Data
- Application Exploitation

ENISA Threat Taxonomy

Physical attack (deliberate/ intentional)	
<ul style="list-style-type: none"> • Fraud • Sabotage • Vandalism • Theft (devices, storage media and documents) • Information leakage/sharing 	<ul style="list-style-type: none"> • Unauthorized physical access / Unauthorized entry to premises • Coercion, extortion or corruption • Damage from the warfare • Terrorists attack
Unintentional damage / loss of information or IT assets	
<ul style="list-style-type: none"> • Information leakage/sharing due to human error • Erroneous use or administration of devices and systems • Using information from an unreliable source • Unintentional change of data in an information system • Inadequate design and planning or improperly adaptation 	<ul style="list-style-type: none"> • Damage caused by a third party • Damages resulting from penetration testing • Loss of information in the cloud • Loss of (integrity of) sensitive information • Loss of devices, storage media and documents • Destruction of records
Disaster (natural, environmental)	
<ul style="list-style-type: none"> • Disaster (natural earthquakes, floods, landslides, tsunamis, heavy rains, heavy snowfalls, heavy winds) • Fire • Pollution, dust, corrosion • Thunder stroke • Water 	<ul style="list-style-type: none"> • Explosion • Dangerous radiation leak • Unfavorable climatic conditions • Major events in the environment • Threats from space / Electromagnetic storm • Wildlife
Failures/ Malfunction	
<ul style="list-style-type: none"> • Failure of devices or systems • Failure or disruption of communication links (communication networks) • Failure or disruption of main supply 	<ul style="list-style-type: none"> • Failure or disruption of service providers (supply chain) • Malfunction of equipment (devices or systems)
Outages	
<ul style="list-style-type: none"> • Loss of resources • Absence of personnel • Strike 	<ul style="list-style-type: none"> • Loss of support services • Internet outage • Network outage
Eavesdropping/ Interception/ Hijacking	
<ul style="list-style-type: none"> • War driving • Intercepting compromising emissions • Interception of information • Interfering radiation • Replay of messages 	<ul style="list-style-type: none"> • Network Reconnaissance, Network traffic manipulation and Information gathering • Man in the middle/ Session hijacking
Nefarious Activity/ Abuse	

- Identity theft (Identity Fraud/ Account)
- Receive of unsolicited E-mail
- Denial of service
- Malicious code/ software/ activity
- Social Engineering
- Abuse of Information Leakage
- Generation and use of rogue certificates
- Manipulation of hardware and software
- Manipulation of information
- Misuse of audit tools
- Misuse of information/ information systems (including mobile apps)
- Unauthorized activities
- Unauthorized installation of software
- Compromising confidential information (data breaches)
- Hoax
- Remote activity (execution)
- Targeted attacks (APTs etc.)
- Failed of bussines process
- Brute force
- Abuse of authorizations

Legal

- Violation of laws or regulations / Breach of legislation
- Failure to meet contractual requirements
- Unauthorized use of IPR protected resources
- Abuse of personal data
- Judiciary decisions/court orders

NIST Risk Assessment Threat Event Taxonomy Exemplary

Adversarial

Perform reconnaissance and gather information

- 5 sub-elements

Craft or create attack tools

- 6 sub-elements

Deliver/insert/install malicious capabilities

- 14 sub-elements

Exploit and compromise

- 17 sub-elements

Conduct an attack

- 21 sub-elements

Achieve results

- 13 sub-elements

Maintain a presence or set of capabilities

- 2 sub-elements

Coordinate a campaign

- 6 sub-elements

Non-Adversarial

- Spill sensitive information
- Mishandling of critical and/or sensitive information by authorized users
- Incorrect privilege settings
- Communications contention
- Unreadable display
- Earthquake
- Fire
- Flood
- Hurricane
- Resource depletion
- Introduction of vulnerabilities into software products
- Disk error
- Pervasive disk error
- Windstorm/tornado

Taxonomy of Operational Cyber Security Risks

Actions of People

Inadvertent

- Mistakes
- Errors
- Omissions

Deliberate

- Fraud
- Sabotage
- Theft
- Vandalism

Inaction

- Skills
- Knowledge
- Guidance
- Availability

Systems and Technology Failures

Hardware

- Capacity
- Performance
- Maintenance
- Obsolescence

Systems

- Design
- Specifications
- Integration
- Complexity

Software

- Compatibility
- Configuration management
- Change control
- Security settings
- Coding practices
- Testing

Failed Internal Processes

Process controls

- Status monitoring
- Metrics
- Periodic review
- Process ownership

Supporting Processes

- Staffing
- Funding
- Training and development
- Procurement

Process design or execution

- Process flow
- Process documentation
- Roles and responsibilities
- Notifications and alerts
- Information flow
- Escalation of issues
- Service level agreements
- Task hand-off

External Events

Disasters

- Weather event
- Fire
- Flood
- Earthquake
- Unrest
- Pandemic

Legal issues

- Regulatory compliance
- Legislation
- Litigation

Business issues

- Supplier failure
- Market conditions
- Economic conditions

Service dependencies

- Utilities
- Emergency services
- Fuel
- Transportation

© 2018 The SANS Institute, Author Retains Full Rights

Steve Launius



Upcoming SANS Training

[Click here to view a list of all SANS Courses](#)

SANS Essentials Australia 2021	Melbourne, AU	Feb 15, 2021 - Feb 20, 2021	Live Event
SANS OnDemand	OnlineUS	Anytime	Self Paced
SANS SelfStudy	Books & MP3s OnlyUS	Anytime	Self Paced