Dakota State University
Beadle Scholar

Masters Theses & Doctoral Dissertations

Spring 3-2018

A Capability-Centric Approach to Cyber Risk Assessment and Mitigation

Thomas H. Llansó
Dakota State University

Follow this and additional works at: https://scholar.dsu.edu/theses

Part of the Other Computer Sciences Commons

Recommended Citation

Llansó, Thomas H., "A Capability-Centric Approach to Cyber Risk Assessment and Mitigation" (2018). *Masters Theses & Doctoral Dissertations*. 339.
https://scholar.dsu.edu/theses/339

This Dissertation is brought to you for free and open access by Beadle Scholar. It has been accepted for inclusion in Masters Theses & Doctoral Dissertations by an authorized administrator of Beadle Scholar. For more information, please contact repository@dsu.edu.

A CAPABILITY-CENTRIC APPROACH TO CYBER RISK ASSESSMENT AND MITIGATION

A dissertation submitted to Dakota State University

in partial fulfillment of the requirements for the degree of

Doctor of Science in Information Systems

March 2018

By

Thomas H. Llansó

Dissertation Committee:

Dr. Cherie Noteboom, Co-Chair

Dr. David Bishop, Co-Chair

Dr. Ashley Podhradsky

Dr. Surendra Sarnikar

DISSERTATION APPROVAL FORM

This dissertation is approved as a credible and independent investigation by a candidate for the Doctor of Science in Information Systems degree and is acceptable for meeting the dissertation requirements for this degree. Acceptance of this dissertation does not imply that the conclusions reached by the candidate are necessarily the conclusions of the major department or university.

Student Name: Thomas H. Llansó

Dissertation Title: A Capability-Centric Approach to Cyber Risk Assessment and Mitigation

Dissertation Co-Chair:______________________________ Date: ______________

Dissertation Co-Chair:______________________________ Date: ______________

Committee member:________________________________ Date: ______________

Committee member:________________________________ Date: ______________

Committee member:________________________________ Date: ______________

Committee member:________________________________ Date: ______________

Acknowledgments

I have many people to acknowledge and sincerely thank:

- My committee: Dr. Cherie Noteboom and Dr. David Bishop, whose careful guidance put me on a solid path forward in my dissertation work. Dr. Ashley Podhradsky, whose domain expertise and enthusiasm for the dissertation topic gave me great encouragement. Dr. Sarnikar, whose expertise in Design Science was of tremendous value; I am grateful for his willingness to serve on my committee despite his transfer to another university.
- The BluGen US Government Sponsor, specifically, Mr. John Garstka, of the Office of the Under Secretary of Defense for Acquisition, Technology and Logistics and Command, Control, And Communication (C3), Cyber, and Business Systems (C3CB). The generous financial support from Mr. Garstka's office allowed BluGen research to go from high-level initial concept to operational research prototype in under twenty-four months.
- My employer, the Johns Hopkins University Applied Physics Laboratory, which encourages and financially supports further graduate study for its employees.
- The entire BluGen team, including Martha McNeil, Dallas Pearson, Brooke Boyd, Dr. Michael Smeltzer, Dr. George Moore, Alyson Grassi, and Courtney Tse. Without the team's energy, talents, and creativity, BluGen would never have become a reality.
- The two EVRA teams, who made time to support the research by participating in the application of the EVRA methodology to the Omega system, thus providing a crucial point of comparison for BluGen.
- Last, but not least, my wife, Rochele, whose support allowed me to pursue this academic journey, a small part of our larger 34-year journey together (so far!)

Abstract

Cyber-enabled systems are increasingly ubiquitous and interconnected, showing up in traditional enterprise settings as well as increasingly diverse contexts, including critical infrastructure, avionics, cars, smartphones, home automation, and medical devices. Meanwhile, the impact of cyber attacks against these systems on our missions, business objectives, and personal lives has never been greater. Despite these stakes, the analysis of cyber risk and mitigations to that risk tends to be a subjective, labor-intensive, and costly endeavor, with results that can be as suspect as they are perishable. We identified the following gaps in those risk results: concerns for (1) their repeatability/reproducibility, (2) the time required to obtain them, and (3) the completeness of the analysis per the degree of attack surface coverage.

In this dissertation, we consider whether it is possible to make progress in addressing these gaps with the introduction of a new artifact called “BluGen.” BluGen is an automated platform for cyber risk assessment that employs a set of new risk analytics together with a highly-structured underlying cyber knowledge management repository.

To help evaluate the hypotheses tied to the gaps identified, we conducted a study comparing BluGen to a cyber risk assessment methodology called EVRA. EVRA is representative of current practice and has been applied extensively over the past eight years to both fielded systems and systems under design. We used Design Science principles in the construction and investigation of BluGen, during which we considered each of the three gaps.

The results of our investigation found support for the hypotheses tied to the gaps that BluGen is designed to address. Specifically, BluGen helps address the first gap by virtue of its methods/analytics executing as deterministic, automated processes. In the same way, BluGen helps address the second gap by producing its results at machine speeds in no worse than quadratic time complexity, seconds in this case. This result compares to the 25 hours that the EVRA team required to perform the same analysis. BluGen helps to address the third gap via its use of an underlying knowledge repository of cyber-related threats, mappings of those threats to cyber assets, and mappings of mitigations to the threats. The results show that manual analysis using EVRA covered about 12% of the attack surface considered by BluGen.

Declaration

I hereby certify that this dissertation constitutes my own product, that where the language of others is set forth, quotation marks so indicate, and that appropriate credit is given where I have used the language, ideas, expressions or writings of another.

I declare that the dissertation describes original work that has not previously been presented for the award of any other degree of any institution.

Signed,

Thomas H. Llansó

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Background of the Problem

Merriam-Webster defines cyber as "of, relating to, or involving computers or computer networks (such as the Internet)" (Merriam-Webster, n.d.). A closely related term, cyberspace, is defined as: “A global domain within the information environment consisting of the interdependent network of information systems infrastructures including the Internet, telecommunications networks, computer systems, and embedded processors and controllers.” (Committee on National Security Systems, 2010). Today cyber is ubiquitous; we interact with it daily via smartphones, tablets, and laptops, but it is also all around us in critical infrastructure, avionics, automobiles, manufacturing robots, and "Internet of Things" (Xia, Yang, Wang, & Vinel, 2012) components, such as medical devices, fitness bracelets, electronic assistants (e.g., Alexa, Siri, Cortana (Heater, 2017)), children's toys, thermostats, and even light bulbs. The software in cyber devices is ever more sophisticated, visualizing protein structures, recognizing faces, translating languages, predicting credit-worthiness, and diagnosing diseases.

While the benefits of applying cyber are significant and growing, so too are the associated risks. Cyber attacks can manifest in many forms, such as identify theft, intellectual property theft, ransomware, and website denial of service. They can be triggered, often anonymously, from great distances, as cyber-enabled devices of all stripes are increasingly interconnected across the globe. Experts especially worry about attacks with societal consequences, such as attacks on voting machines, the electrical power grid, transportation systems, government services, and military systems. Along these lines, adverse events, and adverse cyber events in particular, can lead to high consequence impacts, as illustrated in Figure 1.

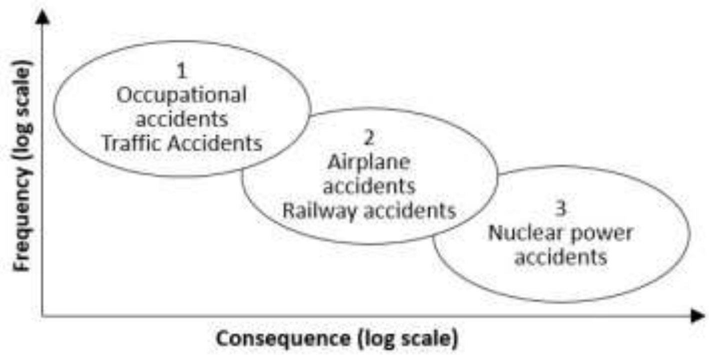

Figure 1: Span of Adverse Events (Rausand, 2011)

The United States Government has grown more concerned about the cyber threat, including within the military, as evidenced by Section 1647 of the 2016 National Defense Authorization Act (NDAA) (Congress, 2016) (Figure 2).

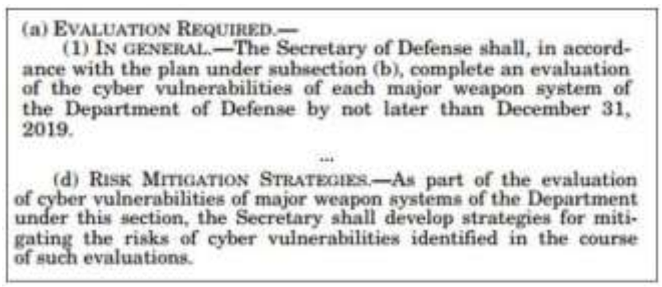
(a) EVALUATION REQUIRED.—
(1) IN GENERAL.—The Secretary of Defense shall, in accordance with the plan under subsection (b), complete an evaluation of the cyber vulnerabilities of each major weapon system of the Department of Defense by not later than December 31, 2019.

...

(d) RISK MITIGATION STRATEGIES.—As part of the evaluation of cyber vulnerabilities of major weapon systems of the Department under this section, the Secretary shall develop strategies for mitigating the risks of cyber vulnerabilities identified in the course of such evaluations.

Figure 2: Excerpt from 2016 Section 1647

Perhaps Congress was motivated by the 2013 Defense Science Board report titled "Resilient Military Systems and the Advanced Cyber Threat" (Gosler & Von Thaer, 2013), which stated:

> *The United States cannot be confident that our critical Information Technology (IT) systems will work under attack from a sophisticated and well-resourced opponent utilizing cyber capabilities in combination with all of their military and intelligence capabilities.*

Regrettably, one might reasonably conclude from the headlines that little has fundamentally changed in the ensuing five years, making the quote as true today as when originally written. Indeed, by nearly any measure, the magnitude of the problem has become staggering. Cybersecurity Ventures estimates that cyber crime will cost the world $6 trillion annually by 2021 and that $1 trillion will be spent globally on cybersecurity from 2017 to 2021 (Cybersecurity Ventures, 2016).

Against this backdrop, organizations that employ cyber systems to help meet their business/mission objectives[1] are concerned about the degree to which cyber attacks can put those objectives at risk. Specifically, with respect to the growing cyber threat, they are interested in answers to a range of questions, such as the following:

- What is my mission risk due to cyber and what mitigations help manage that risk?
- Will the mission survive? Should I limit the use of cyber in the most critical cases?
- As threats, missions, and cyber systems all evolve, how does mission risk change?
- How much risk reduction can be achieved for a given funding level?

Security Architects (SAs) (Newhouse, Keith, Scribner, & Witte, 2017) are on the front-line attempting to help answer such questions. SAs work with other stakeholders, such as managers, mission owners, system owners, other systems engineers, and end users, to make the best decisions possible based on the assessed risk and other considerations, such as funding levels available. SAs typically employ risk assessment methodologies and associated tools to help answer these questions, drawing on others for information required by the assessment process.

A primary output of the risk assessment process is a risk plot, e.g., Figure 3. The plotted data points represent cyber events, such as cyber attacks. Note the ordinal, six-point Likert-style (McLeod, 2008) scale used for each axis in this particular representation. The precise visual depiction of the risk can vary across risk assessment methodologies, but it usually highlights potential cyber events (e.g., attacks) against cyber-enabled components scored by mission impact (also called "criticality" or "consequence;" we use these terms interchangeably) and likelihood of occurrence/probability of success.

[1] Henceforth, this document uses the term ‘mission’ to cover an organization’s business and mission objectives. We note that in government settings, especially the military, the term ‘mission’ is commonly used.

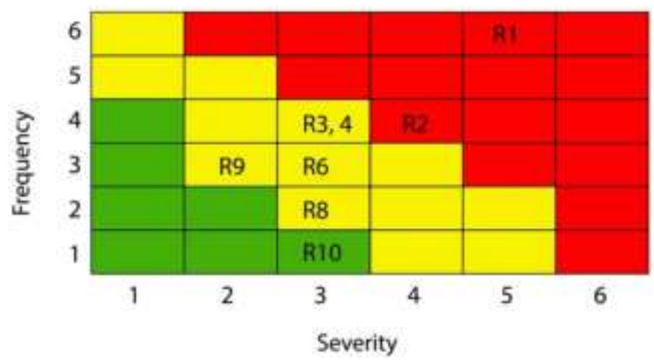

Figure 3: Typical Risk Plot (InsurTech, 2017)

Statement of the Problem

While the SA's are Subject Matter Experts (SMEs) in cyber, their decision-making in risk assessment context is often subjective and variable, leading to concerns about the rigor, repeatability, and reproducibility of the assessed risk and associated mitigation recommendations (Peacos, 2016), (Hallberg, Bengtsson, Hallberg, Karlzén, & Sommestad, 2017). Other concerns include the time and expense required to conduct such assessments. These issues become even more significant given the need to periodically repeat assessments based on the evolution of the (1) anticipated threat, (2) cyber-dependent missions, and (3) supporting cyber systems. In addition, there is growing interest in producing "real time" risk assessment measures for critical systems, making manual assessment unrealistic. Meanwhile, the event-centric approach so commonly employed for cyber risk analysis today has limitations, as captured in part by Aven (Aven, 2016):

> *Traditional risk assessments are based on causal chains and event analysis, failure reporting and risk assessments, calculating historical data-based probabilities. This approach has strong limitations in analyzing complex systems as they treat the system as being composed of components with linear interactions, using methods like fault trees and event trees, and have mainly a historical failure data perspective.*

An additional concern is the need to systematically and objectively identify mitigations that, if implemented, would reduce risk to an acceptable level. Mitigation analysis that is informed by assessed risk and tolerance to that risk is commonly included in the risk evaluation treatment phases of risk analysis. Similar to the scoring of risk, mitigation analysis is typically conducted manually.

Taken together, the concerns discussed above define a gap that the cybersecurity community has historically struggled to address. We return to and expand on these themes in the Literature Review section below.

Research Question

The research question that we pose in this document is as follows: Is there a new approach to mission-cyber risk assessment that can significantly close the following gaps: improved repeatability and reproducibility of results ("repeatability/reproducibility gap"), improved coverage of the attack surface analyzed ("coverage gap"), and decreased analyst time required ("time gap")?

Objective of the project

The objective of the project is to determine the extent to which the gaps mentioned above are addressed by a new approach to assessing mission risk due to cyber effects called “BluGen” (Llanso, McNeil, Pearson, & Moore, 2017)(McNeil, Llanso, & Pearson, 2018). Specifically, the project assesses the degree to which BluGen provides greater coverage of the attack surface and requires less overall SA time to execute for a target cyber system to be analyzed. These time and coverage results are compared to the same results for a representative "first generation" manually-executed, event-centric risk assessment methodology. The project deliverables consist of coverage comparisons, timing comparisons, and an analysis of the extent to which the results support the hypotheses.

CHAPTER 2

LITERATURE REVIEW

We begin the literature review with a basic definition of ‘risk’ and then move on to discuss risk assessment methodologies. The methodology section covers four major categories: compliance-centric, event-centric, loss-centric, and capability-centric. Finally, the review discusses related, cross-cutting topics relevant to cyber risk assessment: mitigation analysis, vulnerability enumeration, human variability in expert scoring, and knowledge management.

Definition of Risk

The assessment and management of risk have been studied for many decades and for many domains beyond cyber, including finance, insurance, healthcare, and military domains including kinetic attack, radiation, and electromagnetic jamming. Despite this long history, there remains a lack of consensus on a single definition of risk. As Kaplan stated in 1997 (Kaplan, 1997):

> *“Many of you remember that when our Society for Risk Analysis was brand new, one of the first things it did was to establish a committee to define the word 'risk.' This committee labored for 4 years and then gave up, saying in its final report, that maybe it's better not to define risk. Let each author define it in his own way, only please each should explain clearly what way that is.”*

Consistent with the quote above, we find many risk definitions in use (Table 1). We note, however, that the definitions all have in common a degree of likelihood or uncertainty with respect to potentially adverse events.

Table 1: Examples of Risk Definitions

Source	Definition
NIST	"A measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence." (National Institute of Standards and Technology, 2012)
ISO 3100	"Effect of uncertainty on objectives" (International Standards Organization, 2009)
Merriam Webster	"Possibility of loss or injury" ("Merriam-Webster Online Dictionary," n.d.)
Investopedia	"The chance that an investment's actual return will be different than expected." (Investopedia Staff, n.d.)
Society for Risk Analysis	"Possibility of an unfortunate occurrence" (Various, n.d.)

The seminal 1981 paper by Kaplan and Garrick, "On the Quantitative Definition of Risk" (Kaplan & Garrick, 1981) captured the essence of these definitions in a more formal way, as follows:

$$Risk = \{< s_i, p_i, x_i >\}$$

In this definition, risk is a set of N events, where an event is represented as a 3-tuple, $< s_i, p_i, x_i >, 1 \leq i \leq N$. s_i is a scenario (event/attack), p_i is the probability of s_i occurring over some defined period of time, and x_i is the consequence (impact) of s_i occurring.

Risk Methodologies

Many existing cyber-related risk methodologies implicitly or explicitly define risk in a manner consistent with the risk definition above, which we call event-centric. In addition to event-centric methodologies, we define three other categories of risk-related methodologies: compliance-centric, loss-centric, and capability-centric. Below, we discuss each of these categories, which are not completely orthogonal from one another, and we provide representative examples of each.

Compliance-Centric Risk Methods

Compliance-centric risk methodologies help organizations comply with policies, such as the Federal Information Security Management Act (FISMA) (House Government Reform

Committee, 2002) and the Department of Defense Instruction (DoDI) 8510.01 (US Department of Defense, 2014). One such methodology that directly supports compliance is the National Institute of Standards and Technology (NIST) Special Publication 800-37, titled "Guide for Applying the Risk Management Framework [RMF] to Federal Information Systems: a Security Life Cycle Approach" (National Institute of Standards and Technology, 2010).

When applying RMF, one undertakes six major steps: (1) categorize an information system, (2) select security controls, (3) implement security controls, (4) assess security controls, (5) authorize the information system, and (6) monitor security controls. Risk is analyzed by considering mission impacts of cyber events in step (1). In step (1), the RMF references two documents to assist in Information System categorization: The Federal Information Processing System Publication 199, "Standards for Security Categorization of Federal Information and Information System" (National Institute of Standards and Technology (NIST), 2004) and the NIST Special Publication 800-60, "Guide for Mapping Types of Information and Information Systems to Security Categories" (Stine, Kissel, Barker, Fahlsing, & Gulick, 2008). Also, in step (1), one analyzes the potential loss of confidentiality, integrity, and availability (C/I/A) of information of various types in a target system and rates the corresponding mission/business impacts due to such a loss along an ordinal scale of Low, Moderate, and High. One then takes the high-water mark rating across all information types as the overall system categorization for the particular loss of C, I, or A. DoDI 8510.01 adopts the NIST RMF, but makes modifications, such as the requirement to use the Committee for National Security Systems (CNSS) Instruction 1253, "Security Categorization and Control Selection for National Security Systems" (*CNSS Instruction No. 1253 - Security Categorization and Control Selection for National Security Systems, Version 2*, 2012). CNSS is similar in concept to FIPS-199.

Discussion. Compliance-centric risk-related methodologies tend to treat risk at a high level. For example, CNSS-1253 considers risk in terms of mission impact/criticality only without regard to the fact that impacts resulting from a compromise of C/I/A can vary for the same mission information across different components of the system and at different times in a given mission time-line. Distinct components might benefit from different mitigation strategies, but the analysis is too high level to differentiate. In addition, CNSS-1253 selects mitigations via lookup tables based on a high-level categorization process. If used alone without a deeper consideration of the full range of risk elements (e.g., threat capabilities, missions, system components, defense capabilities, and mappings among them), one may end up unwittingly over-protecting less important assets, under-protecting more important assets, potentially wasting funds and subsequently imperiling missions. Similar compliance approaches are in use in non-government settings. For example, the Payment Card Industry Security Standards Council (Orfei, Leach, King, Mauro, & Fitzsimmons, 2006), have likewise encouraged a compliance-oriented approach to security with their PCI-DSS security standard.

Event-Centric Methods

Event-centric methods analyze risk by enumerating potential cyber events, such as malicious cyber attacks, and scoring risk as a function of (a) mission impact/criticality and (b) likelihood of occurrence or estimated level of effort to carry out. An event can be malicious (e.g., cyber attack) or non-malicious (e.g., operator error, software error, an earthquake that knocks out electrical power to cyber components). Perhaps the most prominent example of an event-centric methodology is NIST Special Publication 800-30, "Guide for Conducting Risk Assessments" (National Institute of Standards and Technology, 2012), summarized in Figure 4.

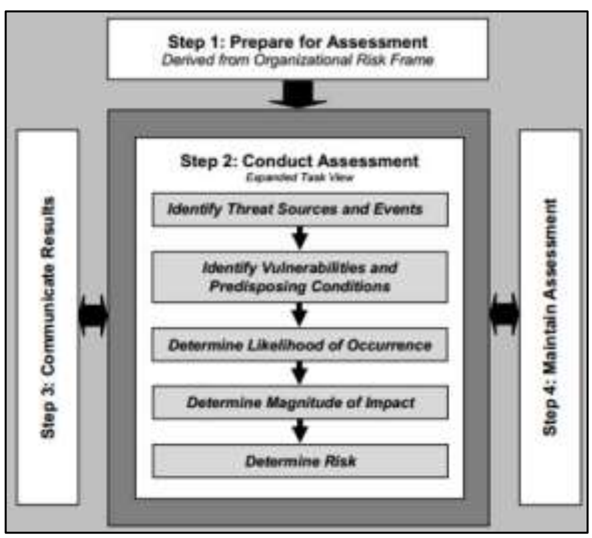

Figure 4: NIST 800-30 Risk Assessment Framework

Other examples of event-centric approaches include the International Standards Organization 27001 Risk Analysis process (*ISO/IEC 27001:2013 - Information technology, Security techniques, Information security Management systems, Requirements*, 2013), Factor Analysis of Information Risk (FAIR) (Carlson, Hutton, & Gilliam, 2010), and the Carnegie Mellon Software Engineering Institute's Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) methodology (Caralli, Stevens, Young, & Wilson, 2007).

Discussion. There are several challenges with conducting event-centric assessments, as typically practiced today. Such approaches usually require mission and SAs to manually score mission impact/criticality and attack likelihood, respectively. However, manual scoring does not scale well due to the combinatoric explosion that results when one attempts to enumerate all possible attack sequences that could be applied to a target system. For example, attack-based risk analysis of a system of 5 mission threads, 40 nodes (computing devices of various types), 3 data items per node on average, 4 attack vectors, and 3 attack effects can require an upper bound of $7{,}200$ $(5 \times 40 \times 3 \times 4 \times 3)$ unique attack contexts that SAs must score for impact and likelihood.

As a result of this combinatoric explosion, SAs tend to consider just a portion of the attack surface by using small, commonly non-random samples, with attendant concerns about how well such samples generalize to the entire attack surface. The result is limited attack surface coverage. Also, such assessments are time consuming and subject to the effects of SA-bias in assigning scores along ordinal scales. Furthermore, the repeatability and reproducibility of such analyses are a concern. While modest progress has been made in automating impact scoring, e.g., (Musman, Tanner, Temin, Elsaesser, & Loren, 2011) and (Llanso & Klatt, 2014), approaches to automating full attack likelihood scoring remain in their infancy. Lastly, to-date there is no clear-cut automation path that leads from attack-centric risk assessment to mitigation analysis, though some related work is going on in this area (Vigo, Nielson, & Nielson, 2014).

Loss-Centric Risk Methods

Loss-centric methodologies are similar to event-centric methodologies described above but are more focused on quantifying dollar losses due to cyber events rather than on assessing potential mission impacts. Two representative examples of such methodologies include the approach described by Seiersen and Hubbard in their book, "How to Measure Anything in Cybersecurity Risk" (Hubbard & Seiersen, 2016) and INFOSEC Institute's "Quantitative Risk Analysis" method (INFOSEC Institute, 2013). In the latter, one determines potential annualized losses to attacks on assets. The key formula in methods similar to INFOSEC Institute's method is $ALE = SLE \times ARO$, where ALE is Annualized Loss Expectancy, SLE is Single Loss Expectancy, and ARO is Annualized Rate of Occurrence. In turn, $SLE = AV \times EF$, where AV is asset value and EF is exposure factor (percent of asset affected by a cyber attack).

Discussion. While potential dollar loss is certainly a reasonable focus for risk, loss-centric methods that approach risk analysis via event enumeration suffer from the same issues as the more mission-focused event-centric methodologies discussed above. Another challenge with such methods is in accumulating enough data to make credible estimates of, for example, ARO and EF. Finally, such methods do not apply as well in situations, such as national defense, where the focus is less about dollar loss and more about mission success and lives saved.

Capability-Centric Methods

The capability-centric approach represents a recent departure from the more common event-centric risk approaches. The idea is as follows: rather than attempting to enumerate and analyze all of the attacks that an adversary might compose from their list of offensive capabilities, the analyst instead focuses on the base capabilities themselves. For each offensive capability, the analyst identifies potential defensive capabilities that could effectively mitigate the offensive capability. Examples of offensive and defensive capabilities at different abstraction levels are given in Table 2.

Table 2: Examples of Offensive and Defensive Capabilities

Level	Example Offensive Capability	Example Related Defensive Capability
1	Threaten system availability	Defend system availability
2	Inject stealthy software implants	Detect and block most stealthy implants via software whitelisting
3	Software implants are injected via air gap jumping methods	Establish an authoritative repository of cryptographic hashes of authorized software

Two example approaches that employ the capability-centric approach are BluGen (Llanso et al., 2017), the focus of this dissertation proposal, and the capability-based approach employed by the government program called "NIPRNet/SIPRNet Cyber Security Architecture Review" (NSCSAR) (Dinsmore, 2016)[2]. BluGen is discussed in greater detail below. NSCSAR focuses on common infrastructure assets used by many missions and considers the degree of exposure of such assets to the anticipated threat, omitting mission impact/criticality considerations.

Discussion. The central hypothesis of the capability-based approach is as follows: as the individual capabilities possessed by an anticipated adversary are mitigated by cyber defenders using their own "defensive" capabilities, it becomes increasingly difficult for that adversary to compose viable attack sequences, because there are fewer and fewer remaining unmitigated "defensive" capabilities from which to compose such attacks. Of course, implicit in this statement is the ability to enumerate the capabilities of the anticipated adversary in the first place, but we believe that this is a more tractable challenge than, for example, enumerating all possible attacks that one could compose from the base capabilities. The 2015 threat model (DoD, 2015) created for DoD provides an example of capability enumeration for the six cyber attacker tiers defined by Gosler and Von Thaer (Gosler & Von Thaer, 2013).

[2] The program name recently changed from NSCSAR to DoDCAR.

Other Related Work

This subsection covers other related work relevant to cyber risk assessment, specifically mitigation analysis, vulnerability enumeration, human variability in expert scoring, and knowledge management.

Mitigation Analysis

Once risk has been assessed, an important next step in the risk assessment and management realm is risk treatment, which examines potential mitigations (also known as countermeasures or security controls) to help manage risk. Representative examples include:

- Step 2 of the Risk Management Framework (RMF) (NIST, 2010)
- Step 4 of ISO 31000 (International Standards Organization, 2009)
- Step 8 of OCTAVE (Caralli et al., 2007)
- Step 2.1.2 of MITRE's Threat Assessment & Remediation Analysis (TARA) (Wynn, Whitmore, Upton, & Spriggs, 2011))

When looking across these steps, we find that they tend to be conducted manually to one degree or another. For example, CNSS-1253 (*CNSS Instruction No. 1253 - Security Categorization and Control Selection for National Security Systems, Version 2*, 2012), which is a recommended approach for realizing RMF step 2, takes a hybrid approach, where the SA consults a large security control table (Table D-1 of Appendix D, Security Control Tables) and mechanically gathers a list of the mapped security controls specified for given levels of mission impact based on breaches of confidentiality, integrity, and availability. Such mappings can be blunt instruments, requiring further SA analysis. The SA then considers possible application of "overlays" (list of controls recommended for particular circumstances, such as systems that include cross domain solutions or that process classified information). Next the SA revises the list (additions/deletions) based on local needs and maps the controls to applicable assets in the target system. An important part of mitigation analysis is consideration of the larger tradespace of cost vs. benefit. The primary benefit is the degree of risk reduction resulting from mitigation. Cost can include a complex set of factors, such as the cost to acquire, integrate, and operate mitigations. Cost can also include negative impacts to the missions of the system caused by use of the mitigations. An extreme illustration of a negative impact would be applying a screen saver that requires a password to an airliner flight deck display. Clearly, such a mitigation could be disastrous during operations such as landing. Tradespace analysis has received some attention in the literature, including work by Dewri, et al. (Dewri, Poolsappasit, Ray, & Whitley, 2007) and Yevseyeva (Yevseyeva, Basto-Fernandes, Emmerich, & van Moorsel, 2015). Dewri takes a multi-objective optimization approach based on attack tree, whereas Yevseyeva employs ideas from portfolio optimization to select security controls. The BluGen team is in the process of considering potential application of genetic algorithms to help search the tradespace of possible security architectures (no published work yet).

Vulnerability Enumeration

Event-centric and loss-centric risk approaches discussed above depend upon the concept of vulnerability enumeration. By vulnerability enumeration, we mean attempting to identify and analyze all the vulnerabilities in a target system. For example, the “Conduct Assessment” step of NIST 800-30 (National Institute of Standards and Technology, 2012), the Risk Identification step of ISO 31000 (International Standards Organization, 2009), and phase 2 of OCTAVE (Caralli et al., 2007) all attempt some form of vulnerability enumeration. While the idea of vulnerability enumeration appeals to the intuition, we assert that for complex cyber environments, attempts at enumerating vulnerabilities will generally fall well short of the total possible set. Therefore, the majority of events that depend on vulnerability enumeration in target systems will not be identified and the related assessment results will thus be incomplete. Undercounts result from the failure to consider exploitation events tied to so-called "zero day" vulnerabilities, that is, vulnerabilities that are known to only a few or not yet known by anyone. Underlying this viewpoint is the paper "Estimating Software Vulnerability Counts in the Context of Cyber Risk Assessments" (Llanso & McNeil, 2018), which analyzes vulnerability discovery rates and the rate of flaw and related vulnerability introduction during the development cycle. The paper combines the two rates in an equation that estimates the number of unknown vulnerabilities as a percentage of total vulnerabilities. The results are not encouraging, with greater than 50 percent of vulnerabilities remaining latent.

Human Variability in Expert Scoring

A theme running through the event-centric and loss-centric risk assessment methods discussed earlier in this section is the routine use of human experts to enumerate events and then score those events for likelihood of occurrence and mission impact. Using humans for this purpose leads to concerns about repeatability and reproducibility of the results. The phrase "inter-rater reliability" is used in the literature (Trochim & Donnelly, 2008) to refer to this issue. As Trochim states:

> *"Whenever you use humans as a part of your measurement procedure, you have to worry about whether the results you get are reliable or consistent. People are notorious for their inconsistency. We are easily distractible. We get tired of doing repetitive tasks. We daydream. We misinterpret."*

While inter-rater reliability has been studied in general settings (Holm, Sommestad, Ekstedt, & Honeth, 2014), (Bolger & Wright, 1994), we focus here on the risk assessment context. Hallberg and his colleagues (Hallberg et al., 2017) studied inter-rater reliability with respect to humans manually scoring the probability and severity of cyber events or incidents. Their study involved 20 raters who scored 105 cyber incidents. After analyzing the results, the researchers concluded that:

> *"The ratings of probability and severity are not reliable enough between raters to be considered a sound basis for the quantification of information security risks."*

Knowledge Management

The discipline of knowledge management (KM) appears to have great potential in the area of cyber risk assessment. Becerra-Fernández and Sabherwal (Becerra-Fernandez & Sabherwal, 2010) define knowledge in a given area as "justified beliefs about relationships among concepts relevant to that particular area." Those same authors define knowledge management, in turn, as "doing what is needed to get the most out of knowledge resources." Activities include the creating, updating, distributing, and employing of knowledge to help address organizational challenges, or, alternatively, per O'Dell and Hubert (O'Dell & Hubert, 2011), “knowledge management is a systematic effort to enable information and knowledge to grow, flow, and create value.”

We see the beginnings of knowledge management in cybersecurity that is relevant to cyber risk assessment. For example, Llansó (Llanso & Engebretson, 2016) defined a model, a subset of which is shown in Figure 5, that captures the details of and relationships between cyber systems, the missions they support, and the cyber threats to which they are exposed. The model, expressed in the Unified Modeling Language (Object Management Group, 1999), captures cyber-related knowledge in six different segments of a unified model. This model was highly influential in the development of the BluGen Reference Catalog (RefCat) discussed in this dissertation.

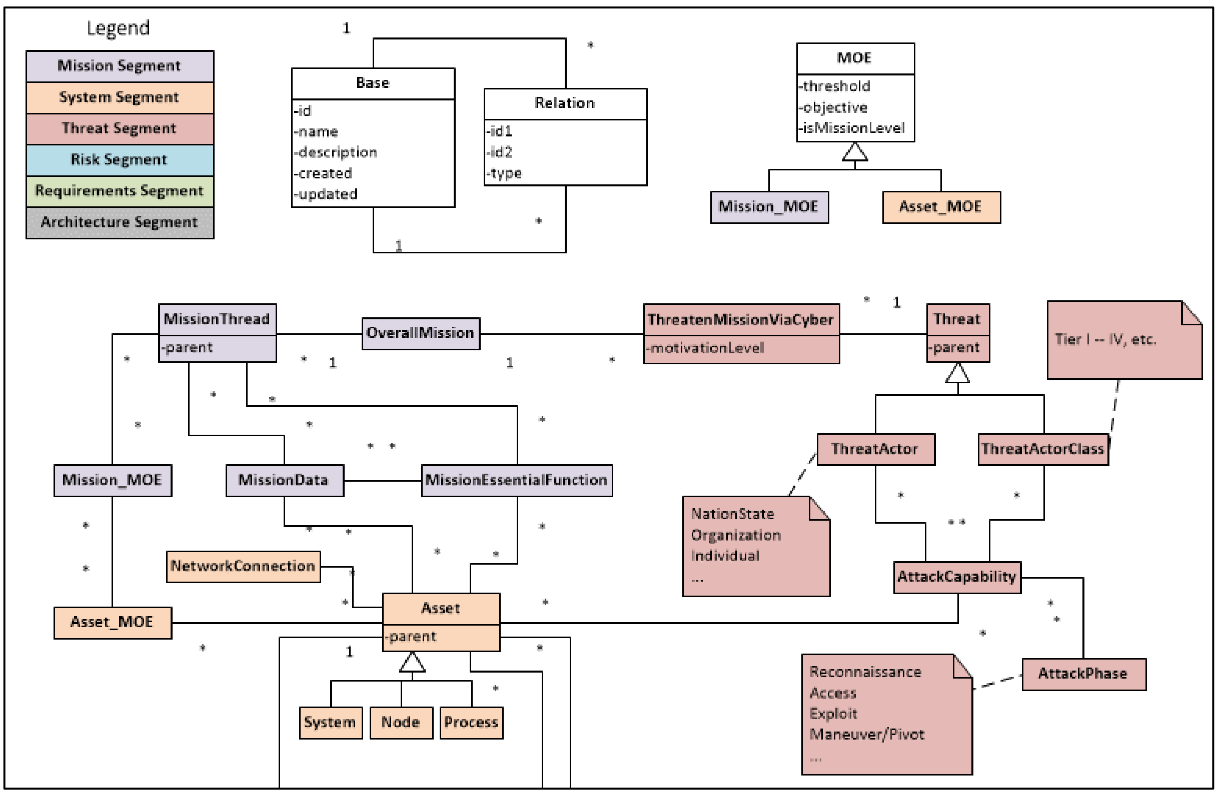

Figure 5: Unified Model for System Security Engineering (UAMSSE) subset

Other cybersecurity models that contribute to the area of knowledge management in cybersecurity include the following:

- D’Amico, Goodall, and Kopylec (Goodall, D’Amico, & Kopylec, 2009) defined a cybersecurity-related model, specifically an ontology that facilitates the mapping of cyber assets to the missions they support and the identification of users who employ the systems composed of those assets.

- NIST's Special Publication 800-53 (*National Institute of Standards and Technology Special Publication 800-53 Revision 4*, 2013) enumerates several hundred security controls intended to be used as mitigations to cyber threats. 800-53 plays a key role in the Risk Management Framework (NIST, 2010).
- MITRE's Common Attack Pattern Enumeration and Classification (CAPEC) (Mitre, n.d.) repository is a rich inventory of cyber attack patterns.
- The National Vulnerability Database (NVD, n.d.) is a highly structured inventory of known vulnerabilities affecting cyber systems.
- MITRE's Mission Assurance Engineering (MAE) (Wynn et al., 2011) model maps mitigations to threats (threats are expressed as TTPs (techniques, tactics, and procedures)) (Figure 6). The model also maps TTPs to asset classes. At a high level, MAE has conceptual similarities to the BluGen RefCat, discussed later.

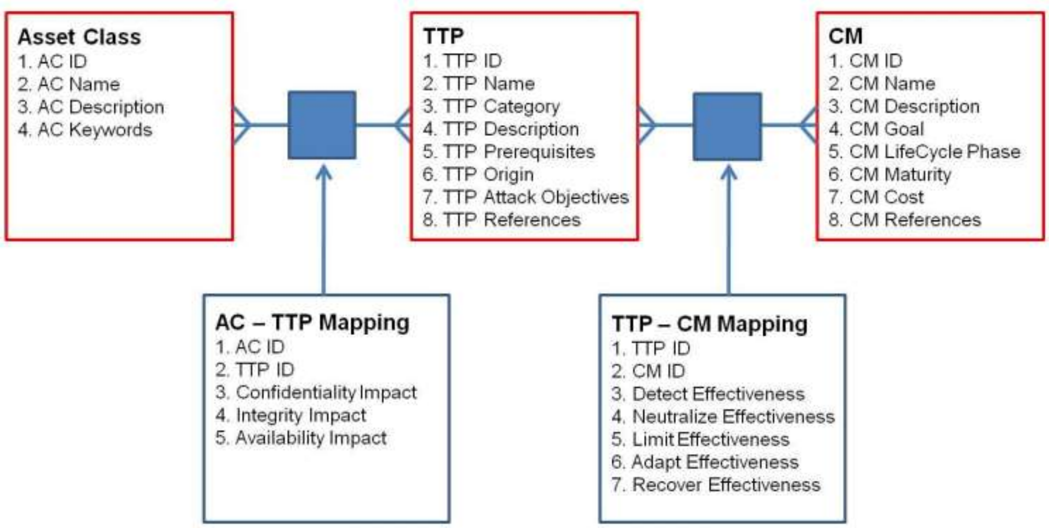

Figure 6: MITRE Mission Assurance Engineering (MAE) Data Model

CHAPTER 3

RESEARCH METHODOLOGY

This section describes the research methodology used, including the placement of BluGen in a Design Science research context, the hypotheses underlying BluGen, the BluGen artifacts themselves, and how we explored those hypotheses. Chapter 4 then presents the results of that exploration.

Design Science Research

Hevner, et al. (Hevner, March, Park, & Ram, 2004) state that “Design science... creates and evaluates IT artifacts intended to solve identified organizational problems.” Vaishnavi and Kuechler (Vaishnavi & Kuechler, 2011) describe Design Science research as “the creation of new knowledge through design of novel or innovative artifacts.” As BluGen consists of a set of designed artifacts, we therefore describe and evaluate BluGen with Design Science Research (DSR) principles in mind.

While different authors approach DSR in different ways, this dissertation adopts the approach described by Peffers, et al., in the 2007 paper titled “A Design Science Research Methodology for Information Systems Research” (Peffers, Tuunanen, Rothenberger, & Chatterjee, 2007). We follow the process in Figure 1, “DSRM Process Model” of that paper, repeated for convenience in Figure 7 below. Our entry point is “Problem-Centered Initiation.” With the problem defined, my research team has been and continues to be in the process of iterating through the steps of that model, which is expected to continue well beyond the timeline of this dissertation. The research behind this dissertation, which has a strong emphasis on the Demonstration and Evaluation phases of Peffers. This dissertation along with other BluGen research already published (Llanso et al., 2017)(McNeil et al., 2018) represents the Communication portion of the Peffer’s Design Science Research Methodology model.

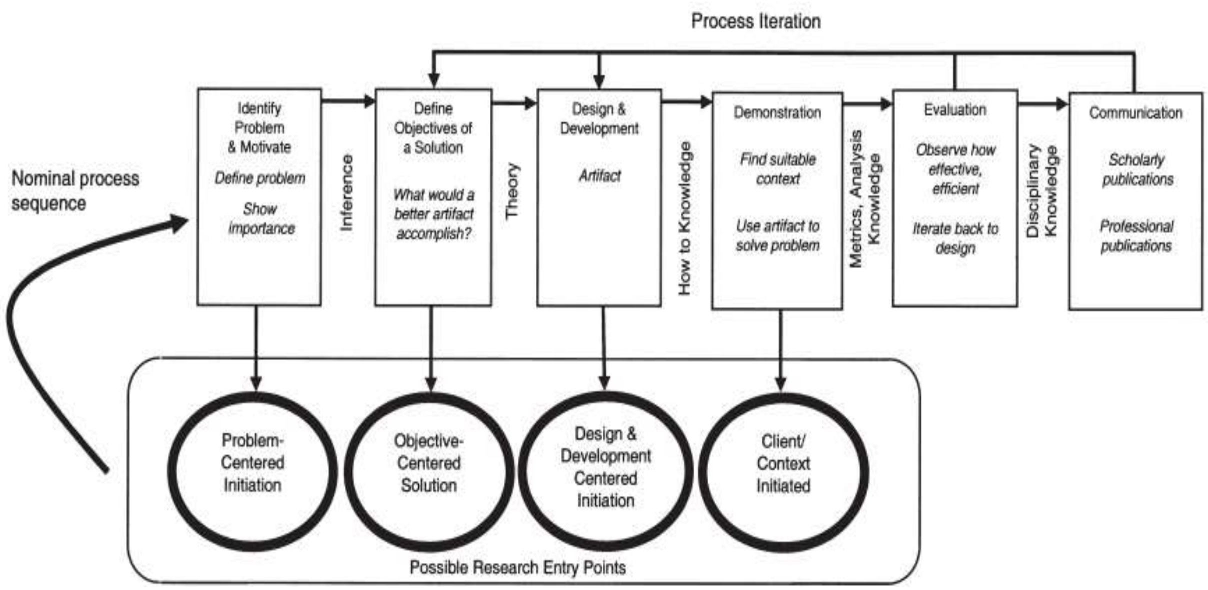

Figure 7: Peffers DSRM Process Model

Theory

As stated earlier, this dissertation centers on BluGen and its evaluation. The hypotheses in Table 3 underlie BluGen. See Figure 8 to place artifacts mentioned in the hypotheses below into an architectural context. The dissertation focuses on hypotheses H1, H2, and H3. The other hypotheses are out of scope and are only included to give the reader a sense of the larger research agenda.

Table 3: BluGen Hypotheses

ID	Hypothesis Summary
H1	BluGen results are more repeatable and reproducible compared to manual, event-centric methods.
H2	BluGen requires less analyst time compared to manual, event-centric methods.
H3	BluGen provides greater attack surface coverage compared to manual, event-centric methods.
H4	Exposure is positively correlated with probability of successful attack.
H5	The following BluGen artifacts have utility to the SA: overall BluGen instantiation, Exposure method, Criticality method, and Mitigation method.
H6	Capability enumeration has utility.

Artifact Design

In this section, we describe the BluGen artifact design. Figure 8 presents a high level architectural view of BluGen. In summary, BluGen is designed as an assistant to the SA and consists of a set of analytic processes and an underlying database called the Reference Catalog (RefCat). To analyze a system for risk and potential mitigations to help manage that risk, the SA prepares a dataset called a "project" that captures essential details about the system to be analyzed and parameters that drive its analysis for risk. The SA submits the project as input to the BluGen software. BluGen analytics cross reference data in the project and RefCat to prepare two major outputs: a risk scatter plot and a report of suggested mitigations (see Table 2 for examples of mitigations).

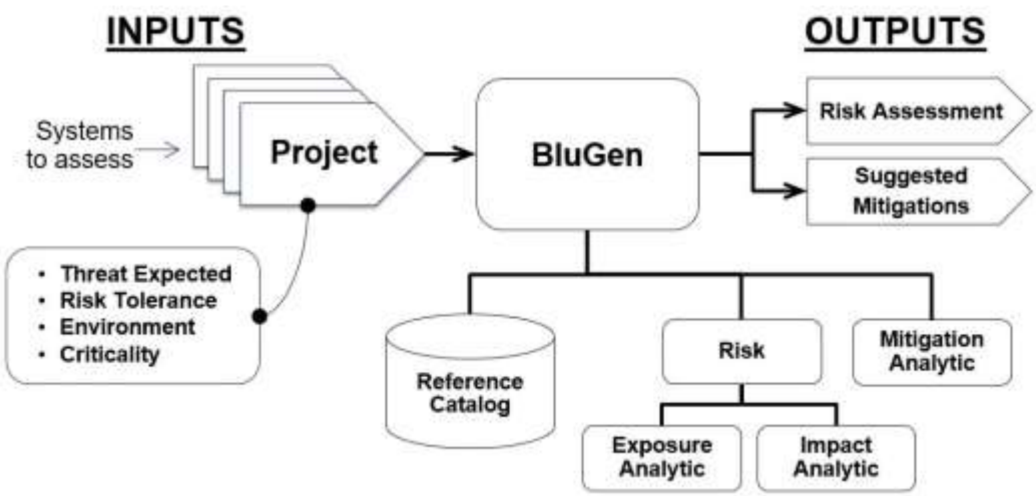

Figure 8: BluGen Architecture

BluGen consists of a number of artifacts, as summarized in Table 4. Below we discuss each of the artifacts using a description adapted and updated from (Llanso et al., 2017).

Table 4: BluGen Artifacts

Artifact	Summary
Framework	BluGen itself
Models	Project and Reference Catalog (RefCat)
Methods	Exposure, Criticality, Mitigation Selection[3]
Instantiation	Java-based instantiation of the BluGen framework

[3] Note that BluGen uses the term 'analytic' to refer to the Design Science concept of 'method'.

Framework

The BluGen framework is the conceptual structure for the capability-based approach for assessing risk and recommending mitigations.

Models

BluGen models consist of the project model and the Reference Catalog model.

Project Model

The project model describes the target cyber environment to be assessed by BluGen and contains three key sets: (1) M, a set of missions; (2) A, a set of assets; and (3) D, a set of data types. The assets in A support the missions in M by processing data in D. Data is subject to compromise possibilities in C, a set of fixed compromise possibilities discussed below. We follow the convention that variables i, j, k, l index objects from M, A, D, and C respectively, under the following four constraints:

- $m_i \in M, 1 \ \leq i \leq |M|$
- $a_j \in A, 1 \ \leq j \leq |A|$
- $d_k \in D, 1 \ \leq k \leq |D|$
- $c_l \in C, 1 \ \leq l \leq |C|$

Each asset instance, $a_j \in A$, consists of a name, an optional description, an asset type, and a set of defensive capabilities that have already been mapped to the asset. The asset type must map onto one of the asset types found in the RefCat model (discussed below). If a new asset type is encountered that is not in the RefCat, it must be added and mapped accordingly. For missions, the environment description includes the overall weight of each mission relative to the other directly supported missions; weights are typically determined by mission and system experts working together. Mission weights should sum to 1.0 for a given Project model instance.

The criticality component of the Project model consists of a set of "raw" criticality 4-tuples. Each criticality 4-tuple, $(m_{i,} a_j, d_k, c_l)$, is a unique combination of four values: a given mission, m, a given asset, a, a given mission data type, d, and a given compromise type, c, chosen from the set $\{CO, IN, AV\}$ where CO represents a breach of confidentiality, IN represents a breach of integrity, and AV represents a breach of availability. Note that not every possible 4-tuple in the Cartesian product of $M \times A \times D \times C$ represents a viable combination, as not every data type is associated with every asset, and not every asset is associated with every mission. Thus, the Cartesian product is an upper bound for the number of tuples required.

Associated with each raw criticality 4-tuple is a score expressed in the range 0.0 to 1.0, with 0.0 meaning not mission-critical at all and 1.0 meaning maximal mission criticality, the worst-case mission impact ("mission kill") if a cyber compromise were to occur in the context defined by the triple. For example, one of many criticality triples for a robot might be: (mission=navigate, asset=sensor, data=location, effect=integrity (IV)) and the worst-case impact for the 4-tuple might be found to be 1.0.

BluGen does not prescribe how raw criticality scores are derived; the scores could be manually assigned by mission experts or they could be generated by a mission/cyber performance simulation that can induce simulated cyber effects and automatically determine related mission impacts, e.g., (Llanso & Klatt, 2014). The former would typically provide scores along an ordinal scale, while the latter would typically provide scores along a ratio scale based on mission performance metrics. The latter is more desirable to help minimize potential SA bias.

Reference Catalog (RefCat) Model

The purpose of the BluGen RefCat model is to capture peer-reviewed cyber- and cybersecurity-related knowledge and make it available for reuse. The BluGen software uses the RefCat along with details about a given target mission/system environment to assess mission risk due to cyber effects (e.g., malicious cyber attacks, human error, acts of nature) and to recommend related mitigations based on a stated threat and risk tolerance. In the realm of knowledge management, the RefCat can be categorized as a knowledge sharing system (Alavi & Leidner, 2001).

The RefCat is a machine-readable repository of cyber knowledge consisting of five primary classes of objects, as follows: (1) a taxonomy of entity types, (2) a set of offensive capabilities that threaten those entity types, (3) a set of defensive solutions that can mitigate offensive capabilities, (4) a set of defensive capabilities from which one composes defensive solutions, and (5) Relationships among the above items. In particular, relationships are many-to-many mappings between offensive capabilities and entity types, defensive solutions and offensive capabilities, and defensive capabilities to defensive solutions.

The RefCat structure is based in part on the model presented in the paper, "A Unified Model for System Security Engineering" (Llanso & Engebretson, 2016). Figure 9 is a summary of the elements above, using a simplified version of Unified Modeling Language notation (Object Management Group, 1999).

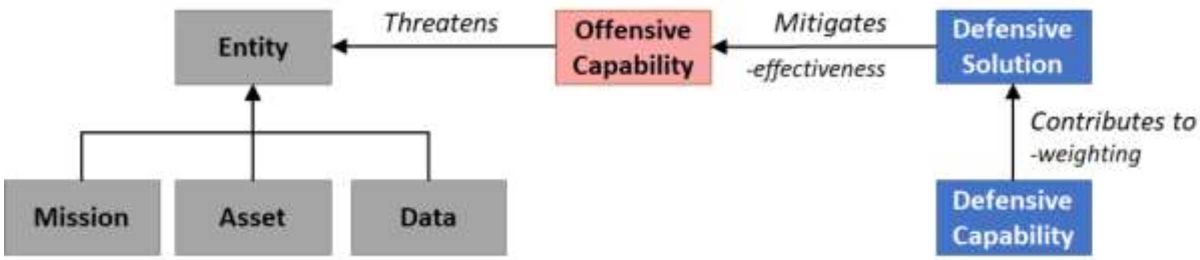

Figure 9: Summary of the RefCat Model (UML)

A few notes on the figure are as follows: Entities can be missions, cyber-enabled assets, and data processed by assets on behalf of missions. Offensive capabilities threaten assets in a many-to-many relationship. Defensive solutions can mitigate offensive capabilities. A defensive solution consists of a set of defensive capabilities mapped to a defensive solution or mapped indirectly via defensive Groups. A defensive group specifies a set of defensive capabilities that are often used together. A defensive model (not shown, but present in the RefCat) consists of a specific set of defensive capabilities that models a particular cyber adversary (e.g., country X, organization Y) or a particular class of adversaries (e.g., Defense Science Board (DSB) tier 3 (Gosler & Von Thaer, 2013)). The RefCat can have many defensive models that represent different subsets of the defensive capabilities recorded in the RefCat. A defensive model consists of a set of defensive solutions and their related defensive capabilities

Methods

This section presents the three major BluGen methods: Exposure, Criticality, and Mitigation.

Exposure Method

In BluGen, we leverage the capability-based representation to define that an entity has higher exposure to anticipated cyber threat actors if it is threatened by a greater number of offensive capabilities for which there are no corresponding set of mitigating defensive capabilities. The Exposure method computes this quantity as presented in Equation 1.

Equation 1: Exposure Method

$$\forall a \in A,\ exposure(a) = \frac{\sum_{oc}^{OC_a} \max(\forall ds \in DS_{rc}\ \sum_{ab}^{AB_{S_{rc}}} (DS_{oc}.weight \cdot dc.weight \cdot present(dc)))}{|OC_a|}$$

Table 5 contains a legend of the symbols used in the exposure and criticality equations.

Table 5: Equation Symbols

Symbol	Meaning
M	Missions in the criticality input data ($m \in M$)
A	Assets in the criticality input data ($a \in A$)
D	Data types in the criticality input data (d ∈ D)
OC_a	Offensive capabilities that threaten an asset type, a (oc ∈ OC)
DS_{oc}	Defensive solution to mitigate defensive capability oc
AB_{oc}	Ability (either solution or defensive group)
$present(dc)$	1 if defensive capability is present for mitigating threat to a; else 0
$mw(m_i)$	Mission weight for the ith mission
$crit(m_i, a_j, d_k)$	Mission criticality score for the given data (d) processed by the given asset (a) on behalf of the given mission (m)
$max(\ldots)$	Entity with the highest criticality
$weight$	Effectiveness of a given solution

As Equation 1 shows, the Exposure method considers each entity in the system, looking up its corresponding entity type in the RefCat. It then searches for all applicable offensive capabilities that are mapped to assets of the given type. Next, for each offensive capability, the method seeks the "best" defensive solution available in the RefCat to mitigate the offensive capabilities, among potentially many solutions available. The best solution is identified by scoring each candidate solution. This is done by summing up the products of the defensive capabilities required for the solution that are present in the current system times the overall solution effectiveness ('weight' in the equation). Lastly, the sum is divided by the number of solutions available to yield a mean effectiveness, which is registered as the overall exposure score for the entity. Figure 9 shows an abstracted example of the exposure analytic.

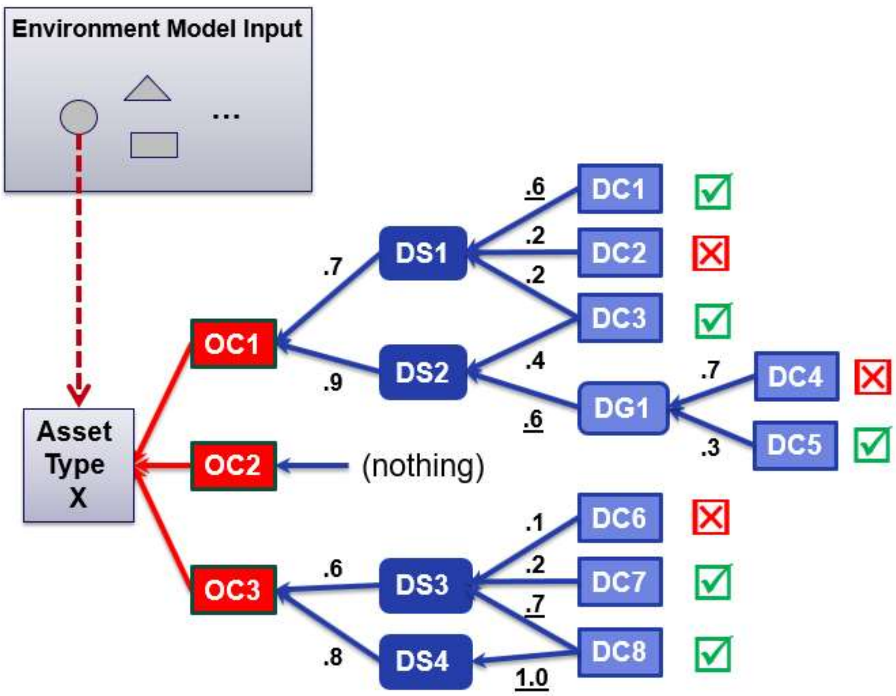

Figure 10: Exposure Analytic Example

In the example (Figure 10), the exposure analytic iterates through the Project model, considering each asset instance in turn. For a given asset, the analytic looks up the corresponding asset in the reference catalog, finding Asset Type X. Next, it looks up the offensive capabilities that threaten assets of that type, finding OC1, OC2, OC3. Then, for each offensive capability, the analytic looks up the defensive solutions that are available to mitigate the offensive capabilities, finding DS1, DS2, DS3, and DS4. Next, the analytic looks up the defensive capabilities that contribute to each of the blue solutions, finding DC1-DC8. The analytic then cross references each defensive capability to see if it is present in the Project model and is mapped to the corresponding asset instance (meaning it is the identified defensive capability contributing to the mitigation of red capabilities that threaten such assets). The green check marks (☑) indicate that the defensive capability-to-asset mapping exists, while the red X's (☒) indicate the mapping is absent. The number on a mapping from a blue solution to an offensive capability represents an estimate of the effectiveness of the blue solution in mitigating the corresponding offensive capability. The number on a mapping from a defensive capability to the corresponding solution represents the weight of the capability's contribution to the overall solution. A number that is underlined means that the capability is required for the solution to be effective at all.

Figure 11 below shows the calculations for the exposure example discussed in Figure 10. The weight of each defensive capability is multiplied by the effectiveness of the overall defensive solution to produce a score. The score is set to zero if any defensive capability required by the solution is missing in the Project model. Summing the scores for each blue solution for each offensive capability results in a coverage score for the blue solution. These are highlighted in yellow in the figure. For each threat, one minus the coverage score produces the exposure. The overall exposure for the asset is the arithmetic mean of the exposure scores for each offensive capability (0.55 in this case).

				In	Carry				
OC1	DS1	Cap	Wt	Sys?	Over	Score			
	0.7	DC1	0.6	y	0.60	0.42			
		DC2	0.2	n	0.00	0.00			
		DC3	0.2	y	0.20	0.14			
						0.56			
OC1	DS2							Score	
	0.9	DC3	0.4	y	0.40		0.40	0.36	
		DG1	0.6	DC4	0.70	n	0.00	0.00	
				DC5	0.30	y	0.30	0.16	
								0.52	
OC2						Score			
						0.00			
OC3	BS3					Score			
	0.6	DC6	0.1	n	0.00	0.00			
		DC7	0.2	y	0.20	0.12			
		DC8	0.7	y	0.70	0.42			
						0.54			
	BS4								
	0.8	DC8	1.0	y	1.00	0.8			
						0.8			
				OC1	OC2	OC3	Sum	# Ocs	AVG
	Covered		Score	0.56	0.00	0.8	1.36	3	0.45
	Exposed		1-Score	0.44	1.00	0.20	1.64	3	**0.55**
			Final Exposure Score						**0.55**

Figure 11: Calculation for the Exposure Example

Criticality Method

In BluGen, an entity is defined as mission-critical if a greater number of highly weighted missions rely on the entity and a greater number of highly critical data types are processed there. The Criticality method computes this quantity as shown in Equation 2.

$$\forall e \in E\ criticality(e) = \frac{\sum_{i=1}^{|M|} \sum_{j=1}^{|D|} mw(m_i) \cdot crit(e,\ m_i,\ d_j)}{\max(rc(e))}$$

Equation 2: Criticality Method

The criticality of a given asset is the sum of raw criticalities in the Environment that are processed by that asset, scaled by the weights associated with the missions that depend on the asset. The final criticality of an asset is expressed as a ratio of the highest criticality of any asset in the target Environment, thus all Environments will have at least one asset with value 1.0. An abstracted example of the criticality analytic is given in Table 6.

Table 6: Criticality Analytic Example

Mission			A1		A2		A3		A4	
Identifier	Weight	Data	Crit.	Weighted	Crit.	Weighted	Crit.	Weighted	Crit.	Weighted
M1	0.3	D1			0.5	0.15			1.0	0.30
		D2			1.0	0.30	1.0	0.30	1.0	0.30
		D3					1.0	0.30		
M2	0.2	D1	1.0	0.20	1.0	0.20			1.0	0.30
		D2	0.5	0.10			0.5	0.10	1.0	0.30
		D3					0.5	0.10		
M3	0.4	D1					1.0	0.40		
		D2	1.0	0.40			1.0	0.40		
		D3	1.0	0.40						
M4	0.1	D1								
		D2	1.0	0.10						
		D3								
Raw Criticality				1.2		0.65		1.6		1.2
Asset Criticality				0.75		0.41		1.00		0.75

Table 6 is mission criticality data from a target environment description provided as part of the input project supplied to BluGen. In this simple example, there are four missions, M1-M4, each with a corresponding mission weight. Mission weight indicates the relative importance of a given mission compared to other missions supported; BluGen expects the weights sum to 1.0. The environment processes three data types, D1-D3, and the data for each mission is mapped to each of the four asset instances (A1-A4). Computing the overall criticality of each asset involves summing up the weighted criticality of each data type processed by each asset for each mission, where such a mapping exists. The sum of the results is then computed, resulting in 'raw criticality' for each asset. The final overall asset criticality is simply the ratio of each raw criticality to the highest raw criticality among the assets considered. Thus, one asset will always have a criticality of 100% using this method (asset A3 in this example).

Mitigation Method

The Mitigation Method, which is a logical extension of the exposure method, recommends mitigations that are currently missing in the target Environment based on the anticipated threat. For each entity in the Environment, the mitigation method looks up the corresponding entity type in the RefCat. Then, for the given entity type, the mitigation method looks up the offensive capabilities possessed by the anticipated adversary that threaten entities of the given type. For each of the offensive capabilities identified, the mitigation method then looks up candidate defensive solutions that map to the given defensive capability. Solutions are assigned a given level of effectiveness, expressed as a percentage, with respect to a given defensive capability. The mitigation method selects the most effective solution and reports a list of the defensive capabilities associated with that solution that are not already implemented in the target Environment.

Instantiation

We created an initial instantiation, Version 1.0, of the BluGen framework implemented in the Java programming language with file storage in Java Script Object Notation (JSON) files. The JSON files contain the RefCat and the environmental/project data.

Exploring the Hypotheses

To test our hypotheses, we carried out a comparative study involving the risk analysis of a ground system for a geosynchronous satellite. For ease of reference, we refer to the ground system as "Omega." The study evaluates the hypotheses by comparing analysis results from BluGen and a representative, manually scored event-centric method. We refer to the event-centric method as EVRA, short for Event-based Risk Analysis[4].

Description of Target System to be Analyzed

Working with a team of experts in aerospace systems engineering, we prepared a detailed description of a ground system that controls a geosynchronous satellite and its payloads. Omega was created as part of an earlier research project. The overall mission of Omega is Space Situational Awareness (SSA). The SSA mission has, in turn, two sub-missions: (1) optical sensing of objects in space and (2) communications of SSA data to various parties. To keep the example openly publishable, the ground system design is a composite of many real ground systems, but the description is not specific to any single ground system. An overview of the ground system architecture appears in Figure 12. The ground system consists of many interconnected cyber components, including controller workstations for the satellite itself and each of the two satellite payloads. The payloads on-board are an optical sensor and a communications transponder. In addition to the hardware and software components identified in the figure, the system also consists of a number of roles that people play to control the satellite and its payload as part of carrying out the SSA mission. Examples of roles are the sensor manager, communications manager, and satellite ops (operations) manager.

[4] Analysts have applied EVRA in over twenty studies, covering both concept-level and fielded systems. EVRA includes an automated tool that provides bookkeeping assistance when logging the manually-scored attacks. The tool also generates risk plots based on these scores.

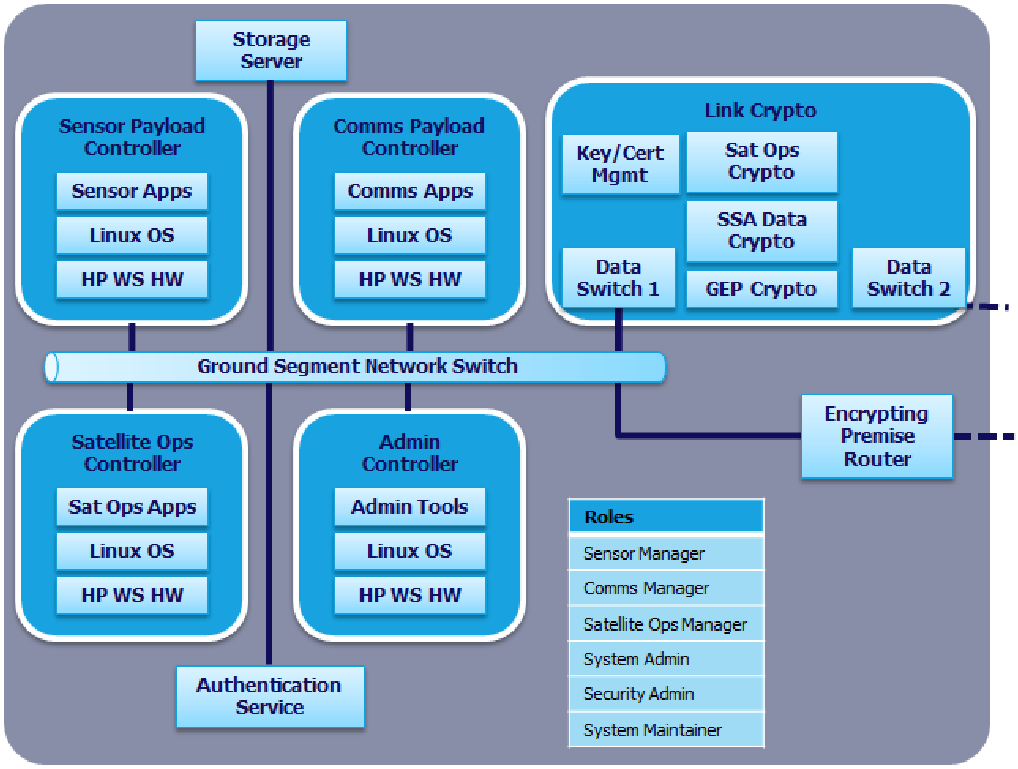

Figure 12: Ground System

Summary information about the architecture in Figure 12 is given in Table 7. As indicated, there are 994 unique entities and relationships in Omega.

Table 7: Entity/Relationship Counts in Omega

Entity/Relationship	Count
Missions	2
Cyber-related asset instances	33
Unique asset types	13
Data types	26
Asset-to-asset mappings (containment)	32
Asset-to-asset mappings (capability inheritance)	80
Data-to-asset mappings	283
Unique existing mitigations	38
Existing mitigations-to-asset mappings	204
Unique Mission-Data Type-Asset combinations	283
Total Entity Count	**994**

Table 8 summarizes the two missions supported by the ground system shown in Figure 12. The information consists of three attributes: a unique identifier (ID column), the name of the missions (Name column) and the relative importance assigned to each mission (Mission Weight column). The relative importance of each mission is represented by a weight value, $0.0 \leq \text{weight} \leq 1.0$, under the constraint that the relative mission weights must sum to 1.0. The weighting information informs EVRA SAs and the BluGen criticality analytic of mission importance when determining the criticality of assets.

Table 8: Missions

ID	Name	Mission Weight
M1	Relay Comms Traffic between SSA Data Customers	0.6
M2	Provide Space Observations to SSA Data Customers	0.4

Figure 13 shows the portion of the RefCat asset type taxonomy relevant to Omega.

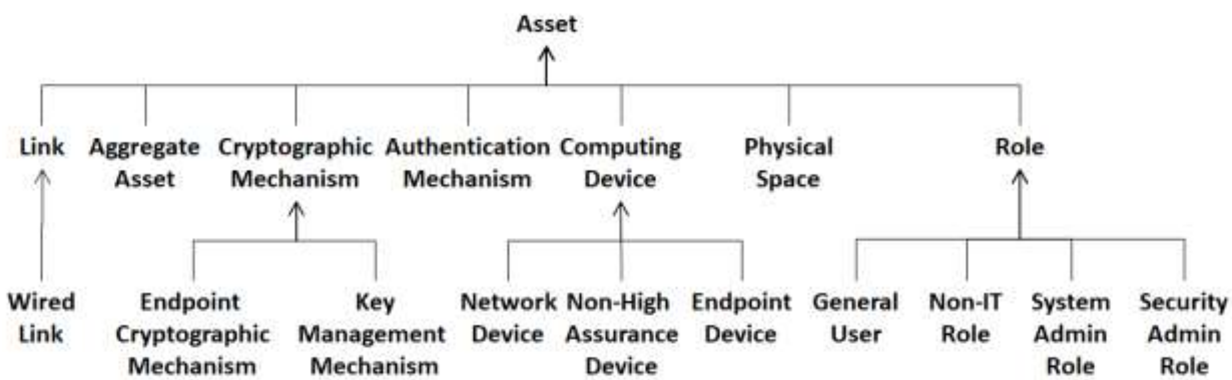

Figure 13: Subset of Asset Type Taxonomy Referenced by Omega

Table 9 augments Figure 13 to include four attributes: (1) a unique identifier (ID), (2) the asset type name, (3) the description of the asset type, and (4) the ID of the parent asset type (PID), if any, for the given asset type. Note that the asset type matches an existing asset type in the asset type taxonomy in the RefCat.

Table 9: Asset Types Descriptions

ID	Asset Type Name	Summary Description	PID
1	Asset	Anything that has value to an organization (other than missions and data), including, but not limited to, another organization, person, computing device, information technology (IT) system, IT network, IT circuit, software (both an installed instance and a physical instance), virtual computing platform (common in cloud and virtualized computing), and related hardware (e.g., locks, cabinets, keyboards).	(nil)
45	Aggregate Asset	An asset that is a container for other assets.	1
5	Endpoint Device	A non-embedded device primarily for use by one user or a group of users (e.g., workstation)	4
4	Computing Device	A machine (real or virtual) for performing calculations automatically (including, but not limited to, computer, servers, routers, switches, etc.). The computing device must be IP addressable (or addressable via an equivalent network protocol).	1
29	Authentication Mechanism	A combination of hardware and/or software designed to authenticate passwords, tokens, biometric and/or other information to identify the user of an account or other resource.	1
22	Link	A communications medium between two communicating computers, without intermediary computing devices. Does not include the exposed interfaces on either end-point.	1
7	Network Device	A non-embedded computing device (other than cross domain solutions) that supports the interconnection of other devices via circuits to form a network; and/or controls/limits the flow of information on that network. As such, network devices are responsible for the layers 3 and/or 4 of the OSI Model.	4
23	Wired-Link	Includes all non-RF links	22
15	Role	All relevant roles that people play with respect to a given target system.	1
16	General User	Includes all persons with any role. The larger the size of this group (and/or the degree to which their allowed access/behavior is less limited), the more likely that the group as a whole has moderate to high impact.	15
41	Non-High Assurance Device	A computing device that is not embedded hardware or high assurance	4

ID	Asset Type Name	Summary Description	PID
19	Non-IT Role	Includes persons in roles where the holders need training regarding physical access, even though they do not access the IT network. This could include cleaners, drivers and others who have information about physical security controls, trash, or who might see information on a whiteboard, desk, or printer. The role holder may be an employee, contractor, a family member, or other confidant.	15
18	Security Admin Role	Includes persons in roles to perform cyber security related functions. These roles are almost always of the highest impact, if compromised.	15
17	System Admin Role	Includes persons in roles to perform system administrative functions (apart from mission specific applications). These roles almost always have high impact.	15
38	Physical Space	A campus, building, floor, suite, room, rack, vehicle, deck, etc. that contains cyber assets.	1
34	Cryptographic Mechanism	A combination or hardware and/or software designed to manage the encryption/decryption of data and control who has the keys necessary to perform these operations.	1
46	Endpoint Cryptographic Mechanism	A cryptographic module, a key storage mechanism, and other parts needed to implement a cryptographic mechanism (except key management) on a computing device	34
35	Key Management Mechanism	A part of a cryptographic module that distributes keys (e.g., directly or through certificates) and limits access to appropriate accounts/person.	34

Table 10 shows the asset instances found in Omega. The table includes references to the types of each asset instance. Asset types are shown above in Figure 13 and Table 9.

Table 10: Assets Instances and Their Types

ID	Name	Asset Type	ID	Name	Asset Type
A01	Ground Control Segment	Aggregate Asset	A18	Premise Router Link	Wired-Link
A02	Type 1 Link Crypto	Aggregate Asset	A19	Data Switch 1 Link	Wired-Link
A03	Admin Controller	Endpoint Device	A20	Comms Manager	General User
A04	Comms Payload Controller	Endpoint Device	A21	Satellite Ops Manager	General User
A05	Satellite Ops Controller	Endpoint Device	A22	Sensor Manager	General User
A06	Sensor Payload Controller	Endpoint Device	A23	System Maintainer	Non-IT Roles
A07	Storage Server	Computing Device	A24	Security Admin	Security Admin Roles

ID	Name	Asset Type	ID	Name	Asset Type
A08	Authentication Service	Authentication Mechanism	A25	System Admin	System Admin Roles
A09	Premise Router	Network Device	A26	Ground Segment Physical Access Control	Physical Space
A10	Ground Control-Ground Entry Point Comms	Wired-Link	A27	Sat Ops Crypto	Endpoint Cryptographic Mech.
A11	Network Link	Wired-Link	A28	SSA Data Crypto	Endpoint Cryptographic Mech.
A12	Ground Segment Network Switch	Network Device	A29	GEP Crypto	Endpoint Cryptographic Mech.
A13	Admin Controller Link	Wired-Link	A30	Data Switch 1	Network Device
A14	Comms Payload Controller Link	Wired-Link	A31	Data Switch 2	Network Device
A15	Satellite Ops Controller Link	Wired-Link	A32	Key/Certificate Management	Key Management Mechanism
A16	Sensor Payload Controller Link	Wired-Link	A33	Storage Server Link	Wired-Link
A17	Authentication Service Link	Wired-Link			

Table 11 lists information about the twenty-six data types processed by assets in Omega. The information consists of two attributes: a unique identifier (ID column) and the names of the data types.

Table 11: Data Types

ID	Name	ID	Name
D01	Captured Observations	D14	Telemetry
D02	Comms Traffic	D15	Repository Data
D03	Processed Observation Data	D16	Spacecraft Operations Plan
D04	Space Vehicle Commands	D17	Tasking Information
D05	Time Slot Assignment	D18	Onboard Clock Adjustment
D06	Access Request for Flight Support Access Node	D19	Telemetry and Command Archive Logs
D07	Communications Configuration Commands	D20	Telemetry Requests
D08	Customer Communications Requests	D21	Authentication Data

ID	Name	ID	Name
D09	Customer Observation Requests	D22	Authorization Data
D10	Sensor Configuration Commands	D23	GEP Control Data
D11	Sensor Observation Schedule Commands	D24	Sat Ops Key Material
D12	Sensor Recalibration Commands	D25	Comms Traffic Key Material
D13	Calibration Data	D26	GEP Control Key Material

The tables above define the missions, asset instances, and data types in Omega. Next, we illustrate various relationship mappings present in the ground system. Table 12 shows a sampling of asset-to-asset mappings, of which there are two kinds: (1) aggregation mappings to show which assets are "contained" within other assets and (2) inheritance relationships to show which assets inherit capabilities associated with other assets.

Table 12: Mapping of Assets to Assets (sampling)

Asset 1	Asset 2	Relationship Type
A01	A02	Contains
A01	A03	Contains
A01	A04	Contains
A02	A27	Contains
A02	A28	Contains
A02	A29	Contains
A03	A08	Inherits Capabilities From
A03	A09	Inherits Capabilities From
A03	A24	Inherits Capabilities From

Table 13 shows a sampling of mappings between data types and assets. In particular, the rows in the table show which assets process the "Captured Observations" data type.

Table 13: Mapping of Data Types to Assets

Data Type	Asset
Captured Observations	Sensor Payload Controller
Captured Observations	Storage Server
Captured Observations	Ground Control-Ground Entry Point Comms
Captured Observations	Ground Segment Network Switch
Captured Observations	Sensor Payload Controller Link

Data Type	Asset
Captured Observations	Data Switch 1 Link
Captured Observations	Sensor Manager
Captured Observations	Storage Server Link

Table 14 shows a sampling of mappings from mitigations to assets. Such mappings indicate to BluGen that a given asset benefits from the corresponding mitigation.

Table 14: Mapping of Mitigations to Assets

Mitigation	Asset
Authenticate All Accounts	Authentication Service
Detect and respond (D&R) to moderately-sophisticated techniques in social settings	Comms Manager
Detect and respond (D&R) to moderately-sophisticated techniques in social settings	Satellite Ops Manager
Detect and respond (D&R) to moderately-sophisticated techniques in social settings	Security Admin
Detect and Respond to Authentication Attacks	Authentication Service
Detect and Respond to comprehensive attacks on Weak Commercial Crypto, Keys managed/stored with Commercial Tools	GEP Crypto

Table 15 shows a sampling of mappings in which a given data type is processed by a given asset in the process of supporting a given mission. For each mapping, mission criticality scores are given for three different situations: (1) a breach of data confidentiality (C column), a breach of data integrity (I column), and a breach of availability (A column). For Omega, a SA manually assigned the scores based on his knowledge of the missions and how the underlying system supports those missions. In general, mission experts provide a written rationale for their mission criticality scores; due to space considerations, we omitted this information.

Table 15: Mission Criticality Mappings

Data Type	Asset	Mission	C	I	A
Comms Traffic	Comms Payload Controller	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Comms Traffic	Ground Control-Ground Entry Point Comms	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Comms Traffic	Ground Segment Network Switch	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6

Data Type	Asset	Mission	C	I	A
Comms Traffic	Comms Payload Controller Link	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Comms Traffic	Data Switch 1 Link	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Comms Traffic	Comms Manager	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Comms Traffic	SSA Data Crypto	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Comms Traffic	Data Switch 1	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Comms Traffic	Data Switch 2	Relay Comms Traffic between SSA Data Customers	0.7	0.6	0.6
Space Vehicle Commands	Satellite Ops Controller	Relay Comms Traffic between SSA Data Customers	0.4	0.8	0.8
Space Vehicle Commands	Ground Control-Ground Entry Point Comms	Relay Comms Traffic between SSA Data Customers	0.4	1.0	1.0
Space Vehicle Commands	Ground Segment Network Switch	Relay Comms Traffic between SSA Data Customers	0.4	0.8	0.8
Space Vehicle Commands	Satellite Ops Controller Link	Relay Comms Traffic between SSA Data Customers	0.4	0.8	0.8

Hypotheses Expectations

Below, we discuss what we expect to find with each of the three hypotheses.

H1 Expectations

We argue that H1 (“BluGen results are more repeatable and reproducible compared to manual, event-centric methods”) is supported with the following justification: Given that BluGen executes a deterministic set of analytics (methods), BluGen will, by definition, produce the same outputs given the same inputs, a result that is independent of the security architect (SA) using BluGen. Thus, we argue for repeatability (the same SA using BluGen at different times but with the same inputs will obtain the same outputs) and reproducibility (BluGen will produce the same outputs given the same inputs regardless of which SA submits the inputs). The utility of this hypothesis is with respect to comparison to manual analysis, where human rater variability tends to be a significant issue. Reliability issues tied to human raters was discussed in the literature review, including concerns about the use of human raters in cyber-related risk assessment (Hallberg et al., 2017). Given the foregoing explanation and justification, we consider that H1 has support and will not discuss it further.

H2 Expectations

Our expectation for H2 (“BluGen requires less analyst time compared to manual, event-centric methods”) is that automated analysis of the type performed by BluGen will execute in a short amount of time (seconds to minutes) compared to the time required to perform similar analysis manually, which experience has shown can take from tens to hundreds of hours depending on target system size and complexity. Thus, BluGen total analysis time is expected to be far shorter than EVRA analysis time when analyzing the same target system. This result would support H2.

H3 Expectations

Our expectation for H3 (“BluGen provides greater attack surface coverage compared to manual, event-centric methods “) is that, on average, BluGen provides greater attack surface coverage than manual event-centric methods. The reasoning is as follows. In EVRA, SAs generally proceed node-by-node[5] in the target system and manually assign a score for the estimated level-of-effort (LOE) required to successfully attack the node. The process of assigning scores is usually based on a team of around three SAs discussing what they know about the nature of each node in question (e.g., its vulnerabilities) and assigning a final score by consensus. In the author's experience witnessing manual scoring sessions tied to different risk methodologies, while SAs may write down a brief rationale for each score they assign, they are not always rigorous during this process. For example, whether due to resource constraints, fatigue, or other reasons, SAs do not always consult and systematically cross reference external sources of information (e.g., threat models, asset taxonomies, vulnerability databases, security control libraries, mappings between these). Given the complexity of modern cyber systems, coupled with the often informal and ad hoc nature of this SA-driven scoring process, we argue that gaps in analyzing the attack surface in terms of threat capabilities are almost certain to occur.

To contrast with the manual process described above, approaches like BluGen automatically consider every possible threat capability known to be possessed by the anticipated threat actor that is mapped to each of the assets that make up a given node. Of course, BluGen is limited to whatever knowledge is currently stored in its RefCat. However, the RefCat is expected to grow in size and accuracy over time, as additional content is added and as peer review and empirical validation of its content proceeds.

Comparative Study Details

As mentioned earlier, we undertook a comparative study to explore the hypotheses. Below we lay out a framework for examining the hypotheses for the comparative study. Next, we discuss the state of BluGen software tool and reference catalog used in the study. We then describe the teams that carried out the respective BluGen and EVRA analyses. Finally, we describe in detail the data submitted as input to each analysis.

---

[5] A node is a computer-type asset in EVRA parlance.

Framework for Examining the Hypotheses

As discussed earlier, hypothesis H1 is considered to be supported and is not examined further. Hypotheses H2 and H3 are fundamentally about comparisons of BluGen to the representative manual, event-centric methodology, EVRA.

Whether using BluGen, EVRA, or some other cyber risk methodology, the high-level steps are generally the same. A brief description of those steps appears in Table 16. Our research examined steps 3, 4, and 5 in the table with respect to hypotheses H2 and H3. The remaining steps (1, 2, and 6) were not considered because the data related to those steps is the same for both analysis methods and is thus considered a constant. The data for steps 1 and 2, in particular, were given identically as input for both the BluGen and EVRA methodologies.

Table 16: Major Cyber Risk Methodology Assessment Steps

Major Assessment Steps	Brief Description
1. Collect and load data	Collect and load data on the anticipated threat, description of the target system, missions supported by the system, and information about risk tolerance.
2. Score mission criticality	Score the mission impact if cyber-related effects (e.g., malicious attacks) occur in the context of every viable combination of mission, asset, and data.
3. Prepare "before" risk plot	Score attack level of effort (EVRA) or exposure (BluGen) for the corresponding attack (EVRA) or asset (BluGen) for the target system as presented.
4. Analyze mitigations	Analyze which potential mitigations might help lower risk to a more acceptable level.
5. Prepare "after" risk plot	Score attack level of effort (EVRA) or exposure (BluGen) for the corresponding attack (EVRA) or asset (BluGen) based on the assumed presence of the mitigations identified in step 4.
6. Prepare and brief report	Prepare a report and associated briefing package of the risk assessment results and associated recommendations to be briefed to appropriate stakeholders.

Table 17 identifies the variables associated with H2 and H3 for assessment steps 3, 4, and 5. The variables for H3 cut across the three assessment steps.

Table 17: Assessment Steps Examined and Their Associated Variables

<table>
<tr><th></th><th colspan="2">Hypotheses and Associated Variables for Data Capture
(variables tracked separate for BluGen and EVRA)</th></tr>
<tr><th>Major Assessment Steps</th><th>H2:
Time</th><th>H3:
Coverage</th></tr>
<tr><td>3. Prepare before risk plot</td><td>T_{PB} - Time to prepare before plot</td><td rowspan="3">C_{AT} - Asset types count
C_{OC} - Offensive capability count
C_{DS} - Defensive solutions count
C_{DC} - Defensive capabilities count
C_{M} - Count of mappings</td></tr>
<tr><td>4. Analyze mitigations</td><td>T_{AM} - Time to analyze mitigations</td></tr>
<tr><td>5. Prepare after risk plot</td><td>T_{PA} - Time to prepare after plot</td></tr>
</table>

Analyzing H2 Data. As the variables in Table 17 imply, to evaluate H2, we tracked the time required by SAs to carry out the analysis for steps 3, 4, and 5 for each approach (BluGen, EVRA). Tracking was done via spreadsheets and a time reporting system. In addition, we extrapolated the time values into the future to address the need for reassessment of the target system. Reassessment is necessary for nearly all systems, as threat, mission, and system all tend to evolve with time, thus limiting the shelf life of earlier assessments.

Analyzing H3 Data. To evaluate H3, we performed a (1) comparison of the H3 counts captured in the table (asset types, offensive and defensive capabilities, mitigations, mappings) and a (2) qualitative comparison of the same data. In both cases, we note and discuss differences. We do these steps separately for BluGen and EVRA. The qualitative comparison considers the relative nature and quality of the data, with special attention paid to potential gaps. As with H2, we discuss future coverage potential based on an evolving RefCat.

Approach is Not Statistical in Nature. As discussed in the dissertation proposal, the quantitative analysis associated with the comparative study that we pursued is not statistical in nature, as one would pursue in formal hypothesis testing. This is because the sample size required to achieve a reasonable margin of error is, for the dissertation, impracticable both in terms our ability to recruit enough qualified SA teams to participate and in funding those SA teams for the time required to execute EVRA studies. For a realistic test, we would need at least three SMEs per EVRA study, and the study lead would need to be experienced in conducting at least one prior EVRA study. In addition, we note that Hallberg (Hallberg et al., 2017) already considered scoring variability in risk assessments at the level of individual raters. Thus, rather than attempting to achieve a statistical result, our analysis is instead a combination of the quantitative aspects (time and count differences between BluGen and EVRA) and the qualitative aspects of the analysis, which examines the differences between the two approaches in the context of the hypotheses examined.

State of BluGen Tool and RefCat

We used version 1.0 of the BluGen software and a snapshot of the RefCat as it existed on June 30, 2017. The state of the BluGen RefCat model used in the comparative study is summarized in Figure 14, which is a screen shot from the BluGen software tool.

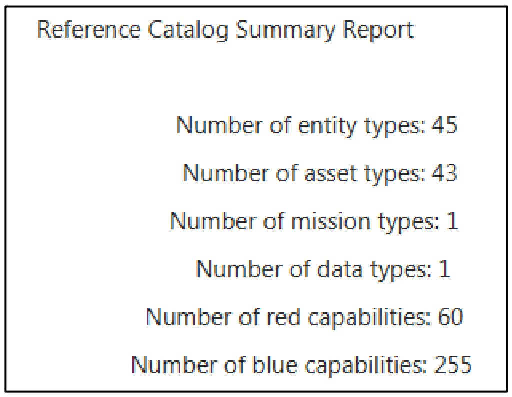
Reference Catalog Summary Report

Number of entity types: 45

Number of asset types: 43

Number of mission types: 1

Number of data types: 1

Number of red capabilities: 60

Number of blue capabilities: 255

Figure 14: Overall Counts in RefCat

The report shown in Figure 14 does not include relationships between capabilities, which numbered 558, and relationships between capabilities and asset types, which numbered 85. Thus, the total number of entities in the Version 1.0 RefCat is 1,048.

Note that both the BluGen software and RefCat continued to be updated iteratively after the Version 1.0 release used for this study.

Hardware Platform for Running BluGen. We ran the BluGen software on a Dell Latitude model E5770 laptop with an Intel Core I7-6820HQ CPU running at 2.7 GHz with 16 GB of main memory and 512 GB of hard disk. On this machine, BluGen was installed as an application on the Windows 7 operating system from Microsoft.

BluGen and EVRA Team Summaries

As the BluGen software conducts the risk assessment and mitigation analysis on a target system automatically, there was no BluGen “team,” per se. The BluGen operator simply instructs the software tool to execute the risk plot generation step and then the mitigations report generation step. As discussed earlier, because we were using BluGen version 1.0, the software lacks the feature to allow the user to check off the desired recommended mitigations based on risk, which is needed to produce the “after” risk plot (the plot produced after accepted mitigations are assumed to be present). A BluGen RefCat developer and a BluGen software developer worked to edit and then reimport the mitigation list. This feature will be automated in BluGen 2.0.

EVRA depends vitally on SMEs for conducting steps such as attack scoring and mitigation determination that BluGen performs automatically. We recruited two separate teams to execute the EVRA methodology for Omega, with the second team acting as a backup to the first team in case the first team was unable to complete the EVRA assessment (e.g., due to personnel availability issues).[6] We used the results from team one to examine H2 and H3. The personnel makeup of both EVRA teams is given in Table 18.

Table 18: Teams That Executed EVRA

Team	Highest Degree(s)	Total Experience
1	BS, Computer Science	6 yrs., 4 mo.
	BS, Physics	3 yrs., 2 mo.
	MS, Info. Technology	0 yrs., 3 mo.
2	MS, Computer Science	6 yrs., 9 mo.
	MS, Computer Science	3 yrs., 5 mo.
	BS, Math; BS, CS	0 yrs., 3 mo.

[6] Having a second team also allowed us to gather anecdotal data concerning the reproducibility aspect of hypothesis H1. We note that the level of scoring consistency between the two teams was poor, with the teams producing different scores for the same attack context greater than 80% of the time.

Inputs to BluGen and EVRA

This section defines inputs to BluGen and EVRA. These inputs are identical except in those cases where there are different input needs between BluGen and EVRA (e.g., the way in which risk tolerance values are described).

Assumed Threat. For our analysis of Omega, we assumed a Tier VI adversary, as defined by the Defense Science Board (DSB) report titled "TASK FORCE REPORT: Resilient Military Systems and the Advanced Cyber Threat" (Gosler & Von Thaer, 2013). Table 19, taken from page 22 and 23 of the report provide brief overview descriptions of the capabilities of the population of threat actors, which can be nation-states, organizations, or individuals, divided among six different tiers.

Table 19: DSB Threat Tier Definitions (Gosler & Von Thaer, 2013)

Tier	Description
I	Practitioners who rely on others to develop the malicious code, delivery mechanisms, and execution strategy (use known exploits).
II	Practitioners with a greater depth of experience, with the ability to develop their own tools (from publicly known vulnerabilities).
III	Practitioners who focus on the discovery and use of unknown malicious code, are adept at installing user and kernel mode root kits10, frequently use data mining tools, target corporate executives and key users (government and industry) for the purpose of stealing personal and corporate data with the expressed purpose of selling the information to other criminal elements.
IV	Criminal or state actors who are organized, highly technical, proficient, well-funded professionals working in teams to discover new vulnerabilities and develop exploits.
V	State actors who create vulnerabilities through an active program to "influence" commercial products and services during design, development or manufacturing, or with the ability to impact products while in the supply chain to enable exploitation of networks and systems of interest.
VI	States with the ability to successfully execute full spectrum (cyber capabilities in combination with all of their military and intelligence capabilities) operations to achieve a specific outcome in political, military, economic, etc. domains and apply at scale.

Referring to the table, Tier I is the least capable threat actor, and tier VI is the most capable. The key assumption underlying the table is that an actor at a given tier n (n > I) possesses the capabilities at the given tier along with all of the capabilities of actors at lower tiers (tiers I through n-1). Thus, for example, a tier III actor possesses the capabilities defined by the union of capabilities across tiers I, II and III. Our assumption of a tier VI threat actor follows from our assertion and that of others (Bateman, 2017) that the most capable nation-states could reasonably have an interest in using cyber as a means to disrupt a system like Omega.

While the DSB report (Gosler & Von Thaer, 2013) defines threat tiers, the tier definitions are defined at too high a level for BluGen analytics or EVRA SAs to conduct their analysis. Both require definition of specific attacker capabilities within each tier. Therefore, to supplement the tier definitions, the BluGen RefCat incorporates a capability definition model that defines capabilities by tier and by category. The model employs seven categories:

- Ability to access networks
- Ability discover and exploit vulnerabilities
- Ability to defeat cryptography and authentication
- Ability affect cyber/physical systems
- Ability to gain physical access
- Sophistication of cyber command and control
- Sophistication of human influence

As an example of a capability, the following is defined for a tier I threat actor in the category called “Ability to defeat cryptography and authentication.” The capability is: “Defeats weak commercial cryptography and weak passwords.” The EVRA SA team was given access to the capability model based on the DSB tiers.

In addition to the threat model mentioned above, EVRA SAs and BluGen RefCat SAs were given the freedom to consider additional capabilities not currently present in the DSB threat capability model.

System. The system description consists of an inventory of assets, including hardware, software, and people (role) assets. For BluGen, we mapped assets instances to their corresponding types in the BluGen RefCat. The description also includes mitigations (defensive capabilities) and various mappings:

- **Connectivity**: which assets connect to other assets via communications links
- **Containment**: which assets contain other assets
- **Mitigation**: which defensive capabilities map to which assets

Other inputs include the following:

- **Mission Criticality**: Mission criticality data, as defined earlier in the section Project Model.
- **Risk Tolerance**: Risk tolerance specifications, which instruct EVRA SAs and the BluGen software as to which assets (BluGen) or attacks (EVRA) are in-scope for active mitigation considerations. Risk tolerance is defined by two variables, as follows:
  - **Mission Criticality.** The mission criticality value on the risk plot above which mitigations are to be considered. For BluGen, mission criticality is on a scale from 0.0 (no mission impact) to 1.0 (complete mission failure). EVRA uses an analogous ordinal scale from 1 to 5.
  - **Likelihood of Impact**. In BluGen, likelihood is estimated via a metric called Exposure, measured on a scale from 0.0 (no unmitigated exposure to the relevant threat capabilities of the anticipated adversary) to 1.0 (full exposure to the relevant threat capabilities of the anticipated adversary). The analogous measure in EVRA is Level of Capability, which is an ordinal scale integer from 1 to 6 to identify the DSB threat tier of the worst-case adversary that possesses the ability to carry out the associated attack event.

CHAPTER 4

RESULTS AND DISCUSSION

This chapter presents the results of the comparative study described in Chapter 3. We begin with a summary of the results for BluGen and EVRA, followed by more detailed results for each. Lastly, we discuss the results in the context of hypotheses H2 and H3.

Summary Results for BluGen and EVRA

Table 20 presents summary results data for BluGen and EVRA for variables defined for hypotheses H2 and H3 in Table 17. Values in the BluGen column for H3 were tabulated from a run of BluGen against Omega, the output of which is summarized in Figure 28 on page 79. We extracted values in the EVRA column from artifacts produced by the EVRA team. See page 87 under the section heading "Omega Data Capture and Timekeeping Data".

Table 20: Summary Data for Hypotheses H2 and H3

Area	Variable		BluGen	EVRA
H2 (Time)	T_{PB} - Time to prepare before plot		<1 sec.	14.30 hrs.
	T_{AM} - Time to analyze mitigations		<1 sec.	5.25 hrs.
	T_{PA} - Time to prepare after plot		12 hrs.	5.40 hrs.
	Totals		**~12 hrs.**	**24.95 hrs.**
H3 (Cover-age)	C_{AT} - Asset types count		13	11
	C_{OC} - Offensive capability count		48	32
	C_{DS} - Defensive solution count		86	N/A
	C_{DC} - Defensive capabilities count		47	16
	C_M - Count of mappings	OffCap→Asset Type	129	45
		DefCap→OffCap	303	N/A
		DefCap→DefSolution	383	N/A
		DefCap→Asset Type	N/A	16
		C_M Total	**815**	**61**
	Totals		**1,009**	**120**

The abbreviations in the mappings portion of the table are: OffCap—Offensive capabilities, DefCap—Defensive Capabilities, and DefSolution—Defensive Solutions. The source of data for the BluGen data is

Table 21 documents key assumptions and characteristics of the Omega analysis, as conducted via BluGen and EVRA. We note that the way in which SAs actually apply EVRA tends to vary from team to team, driven in part by time/funds available and the personality of the team (e.g., whether the team has the patience and endurance to conduct very detailed analysis).

Table 21: Assumptions / Characteristics of the Analyses

#	Assumptions / Characteristics	BluGen	EVRA
1	Considered data types during risk scoring	Yes	No
2	Referred to an explicit threat model	Yes	Yes
3	Maximum assumed threat	Tier VI	Tier VI
4	Mapping of offensive capabilities to asset types	Explicit	Implicit
5	Mapping of defensive capabilities to offensive capabilities	Explicit	Implicit
6	Defensive capability course (RefCat =explicit, SA=implicit)	Explicit	Implicit
7	Analysis includes consideration of different user roles	Yes	No
8	SAs scored EVRA Transit Level of Capability (LOC)	N/A	No
9	Starting nodes (assets) selected in analysis	N/A	All nodes (computers)
10	Attack vectors explicitly considered	N/A	Yes

Explanatory notes on Table 21 are given in Table 22. Values in the # column of Table 22 map back to the corresponding numbered row in Table 20.

Table 22: Explanatory Notes for Table 21

#	Notes
1	The BluGen criticality analytic consults data type information when rolling up criticality scores for 3-tuples of (mission, asset, data-type). Some EVRA teams look at all combinations of mission/asset/data. The team was unable to consider this data due to limited project scope.
2	One advantage that the EVRA team had that teams in the past have generally not had is that we gave the team a copy of a capability-based threat model to work from for LOC scoring. The BluGen threat model in the RefCat is a superset of this model.

#	Notes
3	The DSB Tier VI (Gosler & Von Thaer, 2013) worst-case adversary was used for both BluGen and EVRA.
4/5/6	The use of the term “Implicit” means that the corresponding information came from the heads of the SAs themselves and informal discussions that they had with one another; the consensus results of those discussions were summarized in informal rationale comments recorded by the SA team. The use of the term “Explicit” means that the corresponding information was explicitly recorded in the BluGen RefCat.
7	For time saving reasons, this particular EVRA team chose to ignore multiple user roles and considered one administrative role only.
8	For time saving reasons, this particular EVRA team chose not to explicitly score “transit” LOCs just using target LOCs in their place.
9	The EVRA team considered all nodes in the system as possible starting nodes for attacks. Other EVRA teams sometimes pick just a subset of nodes for this purpose, usually for time savings reasons.
10	Attack vectors (e.g., supply chain, over the network, physical access) are an aspect of offensive capabilities in the BluGen RefCat. In EVRA, SAs have the option of considering them or not as part of the definition of an attack context.

BluGen-Specific Results

We present the BluGen analysis in this section, beginning with screenshots of the BluGen tool after it has been run against the Omega example. The “before” risk plot (meaning *before* any new mitigations are assumed to be applied) appears in Figure 15. The pink shaded region in the upper right-hand portion of the figure is the region of unacceptable risk, which the SA specifies by two input parameters shown at the bottom of the figure: Criticality and Exposure, set in this case to the values 0.50 and 0.25, respectively. These figures taken together mean that any asset instance that has both a criticality score of at least 0.50 and an exposure score of at least 0.25 must be mitigated. The SA considers the risk to the missions from cyber attacks against those assets to be unacceptable. Note that the legend for the assets shown was not fully implemented in this version of BluGen (distinct assets types are supposed to have their own unique icon).

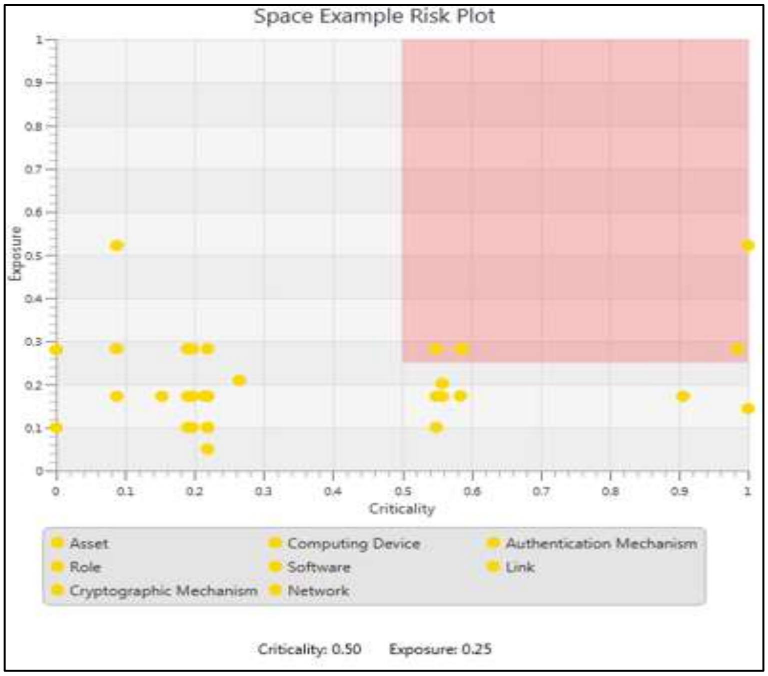

Figure 15: BluGen "Before" Risk Plot

Figure 16 provides a screen shot of the BluGen interactive mitigations report; note the scroll bar on the right. The report extends many pages.

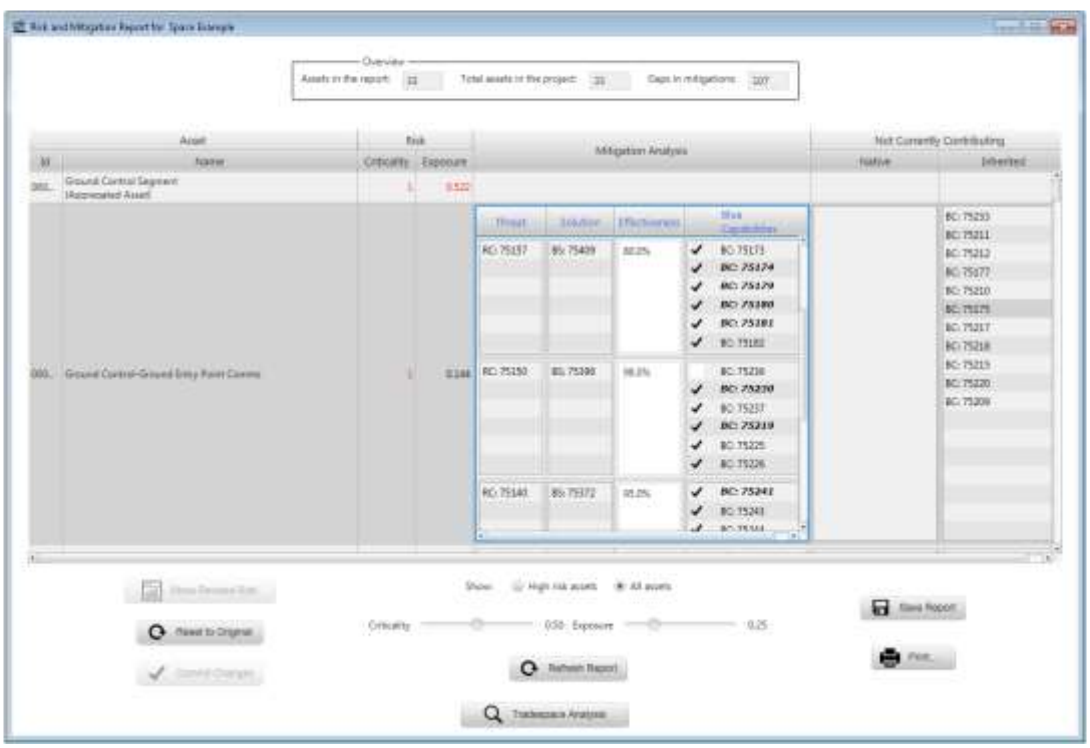

Figure 16: Mitigations Report

The report has one row per asset. The criticality and exposure scores for each asset are shown, with values that exceed the corresponding risk tolerance parameter indicated in red. For each asset, the report shows the red (offensive) capabilities that threaten assets of the corresponding type and the “best” blue (defensive) solution available to mitigate the threat. Also shown are the blue capabilities that make up each solution along with a checkmark that indicates whether the mitigation is currently present or not in the target system.

As mentioned earlier, version 1.0 of the BluGen software does not support the feature, planned for version 2.0, by which a user may selectively choose the mitigations that BluGen recommends for a given target system, threat, and risk tolerance level and instruct the tool to incorporate those mitigations into the model as though they are in place. This feature allows the SA to easily one or more “after” risk plots, show risk under a given set of mitigations. As version 1.0 of the software lacks this feature, the BluGen team manually entered the updated mitigations into the project model and then re-ran the mitigations report. As this was the first time the team had done this, some consultation was required, which took approximately 12 hours total to cover discussions on the best approach, execute the required query, do the manual editing of the mitigations import file, and reimport the file into a revised project.

More information on BluGen data for Omega can be found in Appendix A - Additional Information on BluGen. The appendix includes screen shots and discussion of the BluGen software tool itself as well as special software written to extract the actual coverage data processed by the tool during Omega analysis.

EVRA-Specific Results

We present the EVRA analysis in this section. After the SA’s completed their LOC scoring, they entered those scores along with mission impact scores into the tool so that it could conduct path analysis and generate the risk plot. During path analysis, the software looks at each path from a given starting node in the architecture to a given target node in the architecture, scoring the paths in terms of the SA-provide scores on the LOC for each node along each path. The EVRA tool has no understanding of mitigations, and so does not recommend them, a major difference with BluGen. Instead, the SA’s meet and manually rescore based on mitigations that they devise.

Figure 17 shows the risk plot generated by the EVRA software tool. The number shown by each circle in the plot represents the number of attack contexts that had the same mission impact and LOC scores. The size of the circle is proportional to the number of attack possibilities. The LOC scale is tied to the DSB levels and is thus inverted, so that high-impact, low-capability attacks cluster in the upper right-hand portion of the figure. Color coding emphasis the seriousness of the attacks, with red being the most "risky".

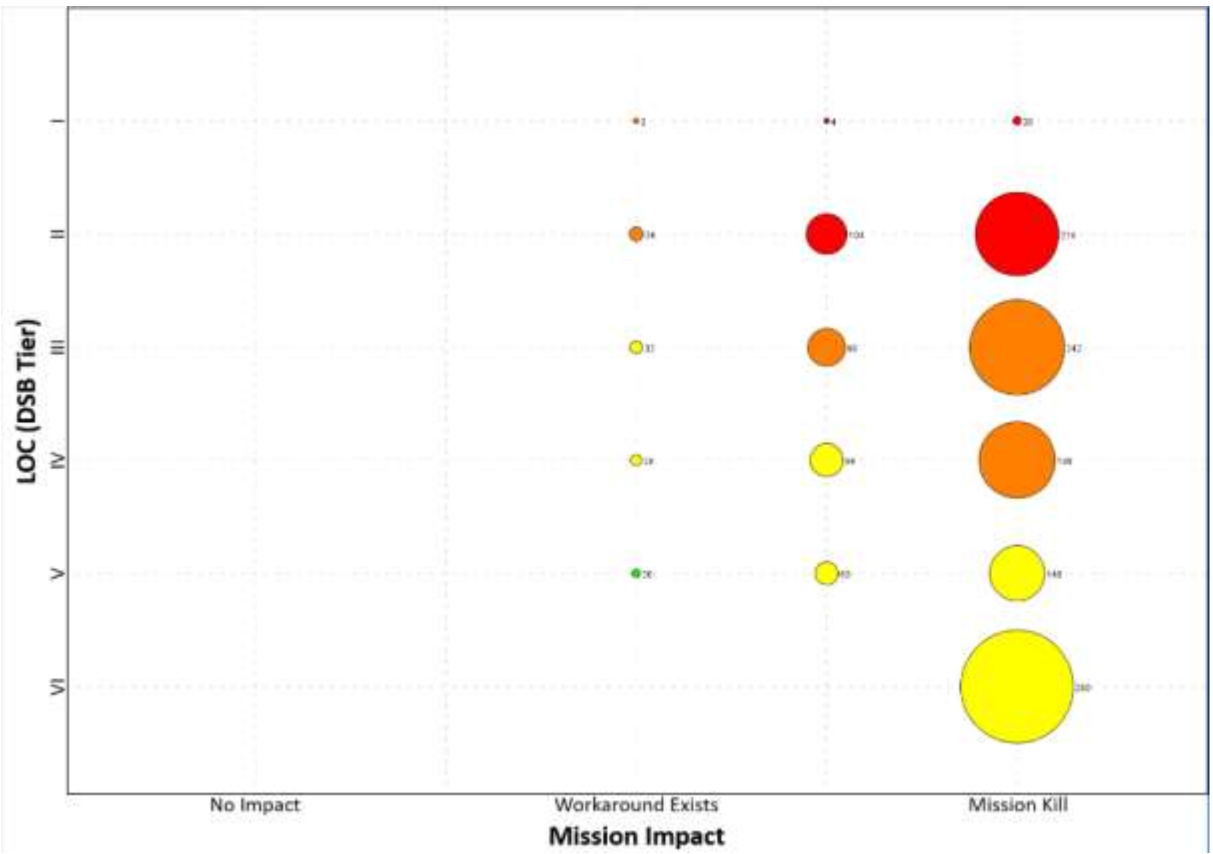

Figure 17: EVRA Risk Plot

More information on EVRA itself as well as Omega scoring artifacts and timekeeping data can be found in Appendix B - Additional Information on EVRA.

Discussion

This section discusses the results in the context of hypotheses H2 and H3. For ease of reference, we repeat the wording of hypotheses H2 and H3 here:

- H2: BluGen requires less analyst time compared to manual, event-centric methods
- H3: BluGen provides greater attack surface coverage compared to manual, event-centric methods.

Hypothesis H2

In the case of BluGen, we found that the three time-related variables, T_{BP}, T_{AM}, and T_{PA}, all took less than one second to execute on the Dell Latitude laptop described earlier. As discussed in the section above, however, BluGen version 1.0 required manual reentry of the mitigation specifications before running the second risk plot, an activity that took 12 hours. In the upcoming 2.0 version of BluGen, this feature will be built into the software, and the user will simply check off the desired mitigations to be incorporated in automated reanalysis. Nonetheless, even considering the time required to manually edit and reimport the external project file, the total time (T_{BP} + T_{AM} + T_{PA}) for the BluGen risk assessment of Omega was still less than half the time required to accomplish the same task by the EVRA team (12 hours vs. 24.95 hours, respectively). In BluGen version 2.0, the time should drop considerably, equating to the time that the SA takes to check a series of boxes indicating whether or not to accept proposed mitigations, which we anticipate to be on the order of a few minutes[7]. Thus, we anticipate support for H2 will grow as further automation comes to BluGen in version 2.0. In consideration of the total time values, we find support for H2.

The total time for an EVRA type analysis is actually magnified by the fact that target systems need to be reevaluated at intervals, such as annually. Reevaluation is needed because the nature of the cyber threat, the mission(s) that a target system supports, and the target system itself, all co-evolve in time, thus limiting the shelf life of any given risk analysis result.

---

[7] This time excludes the time the SA takes to think through implications of selecting different mitigations, which arises whether BluGen, EVRA, or any other risk method is being used.

Ultimately, a desire to assess risk in 'real time' makes the time required to conduct EVRA-style manual analysis untenable.

Hypothesis H3

For H3, the total coverage for BluGen amounted to 1,009 distinct entities vs. 120 for EVRA. Those figures represent totals for the variables C_{AT}, C_{OC}, C_{DC}, and C_{M}, per Table 20. Stated another way, EVRA SAs only considered approximately 12% of the entities compared to BluGen. In consideration of the total coverage values, we find support for H3.

The H3 data for BluGen reflects the state of the RefCat at the time the comparative study was executed. However, as a knowledge repository, RefCat is intended to be under continuous evolution as new cyber asset types are introduced, new offensive capabilities are identified, and new defensive solutions to mitigate the offensive capabilities are designed. In fact, as of this writing (February 2018), the RefCat has grown to 8,953 entities, which is 8.5 times larger than Version 1.0 RefCat used during this dissertation (current as of Summer 2017), which was 1,048 entities. As catalogs such as the RefCat grow in time, the percentage of their content that SAs can reasonably expect to retain "in their heads" so that they can conduct manual risk scoring as they do today is expected to continue dropping. Thus, over the long term, we believe that support for H3 will continue to grow.

In addition to a far richer RefCat, RefCat data quality is expected to improve over time as its contents undergo further peer review and empirical data validation. The idea behind this is that the eventual goal for the RefCat is to host it on servers accessible to the cybersecurity community at large. In this setting, the RefCat will be available not only for reuse but also for peer review of its contents. It is our expectation that data quality will improve through the peer review process, much as academic paper quality can improve when authors take independent reviewer comments into consideration when updating their papers. A level beyond peer review is taking into consideration empirical data from the "real world" cyber environment (e.g., the results of cyber incident response and forensic investigations) and cross referencing that data with data in the RefCat. Assertions in the RefCat can then be squared against the empirical data, acting as another form of quality control. For example, incident data from sensors in major government agencies collected over several months might reveal that the effectiveness of a certain defensive solution recorded in the RefCat is actually lower than the SA-set effectiveness score for the solution in the RefCat (e.g., the effectiveness score might indicate 80% effective, but a large volume of incident data might reveal that the solution is effective only 40% of the time).

Validities

In this section, we consider the validity of the research described above. Valid research is, per Trochim, et al., “the best available approximation to the truth of a given proposition, inference, or conclusion” (Trochim & Donnelly, 2008).

Effort to Create the RefCat and RefCat Sharing with EVRA Team

Before reviewing specific kinds of validities, we first take up a possible point of objection in the manner by which BluGen and EVRA are compared. Specifically, one could argue that while EVRA is a manual method, so too, indirectly, is BluGen, in the sense that the BluGen RefCat is, at least initially, a product of manual (SA) effort. Therefore, an ostensibly fairer comparison of BluGen and EVRA for the time element explored in H2 would have to include the time in BluGen required to manually create the RefCat. Likewise, a seemingly fairer comparison with respect to H3 would involve providing the BluGen RefCat to the SAs for their own reference while executing EVRA. We argue, however, that these concerns are misplaced.

With respect to H2, it is certainly true that the RefCat took time to initially create, and it will likewise take time to maintain and extend the catalog into the future. That said, we expect that this effort will be amortized over hundreds to thousands of automated BluGen analyses that otherwise would have had to have been conducted manually otherwise. In this way, BluGen and its RefCat act as force multipliers.

With respect to H3, while one could provide a copy of the RefCat to SAs as an aid to conduct manual scoring, the goal of BluGen is to replace the need for manual scoring and to provide a means to mitigate the issues that tend to go along with it (reference the prior discussion on this topic and work by Hallberg (Hallberg et al., 2017)). In addition, EVRA represents current practice. We did not want to distort current practice in the context of examining our hypotheses.

Face Validity

A weak form of validity is Face Validity, the extent to which a construct or artifact like BluGen makes sense to others "on the face of it." (Trochim & Donnelly, 2008). We argue for face validity for BluGen in terms of the reactions we have repeatedly experienced when presenting core BluGen concepts to others in the cyber field, specifically, its capability-based nature, its focus on assets, and its particular depiction of risk, including the concept of threat exposure[8]. The approach appears to readily appeal to the intuition of others who are experienced in the cyber risk assessment field.

Instantiation Validity

Lukyanenko, et al. (Lukyanenko, Evermann, & Parsons, 2014) introduced the concept of Instantiation Validity for Design Science Research, which they define as "the extent to which an artifact is a valid instantiation of a theoretical construct or a manifestation of a design principle." They further state that "Instantiation validity is analogous to the concept ... of construct validity in survey research."

We argue for instantiation validity in the sense that the instantiation of BluGen, and in particular, the implementation of the Exposure, Criticality, and Mitigation methods were painstakingly hand-checked by the research team in February and March 2017 against the abstract expression of those methods (equations and corresponding pseudo code). Thus, we have confidence that the instantiation reflects those design concepts.

[8] Since 2016, we have briefed BluGen to a variety of audiences, including HICSS (conference paper), a risk assessment workshop at APL, the International Test Evaluation Association, the US Space Community, and various departments and agencies of the US government.

External Validity

External validity considers whether the results we obtain from our comparative study generalize to other contexts. Below, we discuss the following threats to external validity:

- The target system does not generalize to other system types
- Time results related to h2 do not generalize to larger systems
- EVRA does not generalize to other risk assessment methods
- EVRA team does not generalize to other teams

Threat: Target system does not generalize to other system types. We conducted our comparative study against a single target system, the Omega space ground system. The question is whether we can generalize our results to other target systems that BluGen might be called upon to analyze. Having an insufficient sample size (e.g., a sample of one) would normally be considered a threat to external validity. However, while data from additional investigations conducted against other system types would be welcomed, we do not expect that the results would be materially different in other settings based on the nature of the two hypotheses that we are assessing: time savings and increased coverage.

Threat: Time results related to H2 do not generalize to larger systems. With respect to other system sizes, we have attack and node-related data on eleven previously executed EVRA selected risk assessment studies completed since 2009, as shown in Figure 18.

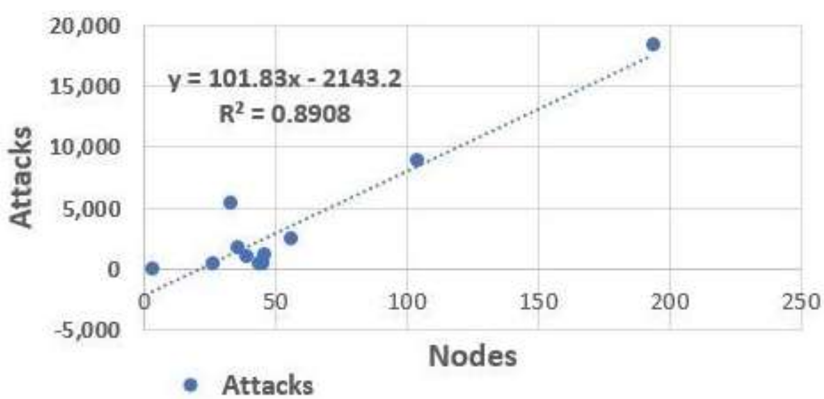

Figure 18: Attacks Analyzed vs. System Node Counts

In this context, the term “node” equates to “computer,” a general kind of asset in BluGen parlance. As the graph shows, the number of attacks chosen by SAs to analyze and score has tended to grow as a roughly linear function of the number of cyber nodes in the target system. As the score for a given attack generally requires a discussion among SAs, the more attacks to be scored, the more total time required to conduct the required discussions. We thus argue that our time results should remain valid across systems of different sizes, as measured by total node count, thus supporting H2 for other systems.

The nature of the BluGen exposure algorithm is on the order of O(n) with worst case O(n^2), where n is the number of assets. The criticality analytic has similar complexity. We base the complexity estimate on the four major nested loops of the exposure algorithm. Below is a simplification of the nested loop procedure:

- Loop A: Consider each asset instance in the environment (n items)
  - Loop B: Consider each offensive capability mapped to the asset's type (m items)
    - Loop C: Consider defensive solution mapped to the offensive capability
      - Loop D: Consider each defensive capability mapped to the defensive solution

We regard the processing time for Loops B, C, and D as equating to a constant factor. On average, we expect that the number of offensive capabilities (Loop B) mapped to an asset instance, m, to be less than 100; the current maximum is 96 and the mean is 34. We expect the number of solutions mapped to an offensive capability (Loop C) to be low (<10) and the number of defensive capabilities mapped to a defensive solution (Loop D) to be even lower (<5) on average. These numbers are based on our experience populating the RefCat thus far. Thus, in summary, the computational complexity of exposure is on the order O(n) or linear complexity.

Threat: EVRA does not generalize to other risk assessment methods. Another possible threat to the external validity is our choice of the comparison methodology, EVRA. If EVRA is not truly representative of attack-based methodologies, against which H2 and H3 comparisons are made, then the argument for external validity is weakened.

However, EVRA conforms to the overall model of the NIST 800-30 Framework (National Institute of Standards and Technology, 2012), which is a commonly accepted approach and a key part of the broadly cited RMF. One notable variance from 800-30 is

EVRA's use of Level of Effort (LOE)[9] in place of likelihood of successful attack on the Y axis. However, we argue that LOE is a legitimate proxy for likelihood in much the same way that we argue that BluGen's exposure method is proxy for likelihood. Other methodologies depend on similar arguments. Indeed, until the community moves away from subjective SA scoring and can collect and analyze sufficient empirical attack data from which to establish frequentist probabilities to support a probability-based Y axis, such arguments are the best we currently have.

To contrast, the analytics that conduct analogous scoring in BluGen operate at machine speeds against data sets that are orders of magnitude smaller than what one would consider to be on the scale of "big data." Thus, we do not expect the set of BluGen algorithms that implement the analytics to encounter a times/space wall for more complex cyber systems than those we have thus far analyzed.

Threat: EVRA team does not generalize to other teams. Table 18 identified the team that executed the EVRA assessment. To evaluate this threat, our main point of comparison is the previously mentioned work of Hallberg (Hallberg et al., 2017). Like the twenty survey respondents in Hallberg's research, the EVRA team members all possess university degrees and have a range of cyber assessment expertise and experience. Hallberg's respondents ranged in age from 29 to 64 years, whereas the EVRA team members are all under 30. A potential limitation to the EVRA team, then, is years of experience, which operates under the premise that additional years of experience correlates to increased expertise for security risk assessment. However, we note that Hallberg concluded the following:

> *"…it cannot be stated that experts have a higher consensus than non-experts when the probability and the severity of information security incidents are rated."*

[9] In some applications of EVRA, such as that described in this dissertation, the SAs score Level of Capability (LOC) rather than Level of Effort (LOE). The former refers to levels of cyber offensive capability in a capability-based threat model, whereas LOE refers to the SA's estimate of "effort" (resources-time/money).

Other Validities: Internal, Construct, Convergent, and Discriminant

H2 and H3 are about comparing selected quantities (time/coverage) between BluGen and EVRA. We argue that internal validity does not apply, as we are not attempting to establish causality in these hypotheses. Likewise, as we are not directly testing a theoretical model in those hypotheses, construct, convergent, and discriminant validities do not apply.

CHAPTER 5

CONCLUSIONS

In 2018, the cyber risk assessment and mitigation process tends to be an SA-intensive effort that is slow, expensive, and has generally poor reproducibility. We again quote Hallberg (Hallberg et al., 2017): "*The ratings of probability and severity are not reliable enough between raters to be considered a sound basis for the quantification of information security risks.*" However, given the ubiquity and critical uses to which cyber is increasingly put, we suggest that the importance of reliable and timely cyber risk assessment results has never been greater. Our original research question was:

> *"Is there a new approach to mission-cyber risk assessment that can significantly close the following gaps associated with what is typically seen in manually executed assessments: improved repeatability and reproducibility of results ("repeatability/reproducibility gap"), improved coverage of the attack surface analyzed ("coverage gap"), and decreased analyst time required ("time gap")?"*

In this dissertation, we introduced BluGen, an automated risk assessment approach that, rather than attempting to enumerate vulnerabilities and possible attack events, focuses instead on underlying attacker capabilities and computes asset exposure to those capabilities along with a rolled-up level of mission consequence. We asserted that BluGen could address the gaps in the research question. To explore whether the evidence supported the assertions, we conducted a comparative study that focused on a target space system, comparing BluGen and a representative attack-centric methodology called EVRA. The basis of comparison centered on three hypotheses tied to the gaps in the research question above (repeatability/reproducibility, time, coverage). Our investigation found support for the hypotheses.

It is our hope that the contribution of BluGen to the knowledge base will help the field of cyber risk assessment and mitigation to become more systematic in its approach and more apt to leverage collected cyber knowledge rather than relying solely on the judgments of individual SAs.

Much work remains to be done. In the context of BluGen, the following elements represent a sampling of areas of possible future work.

(1) **Formal Hypothesis Testing**. To strengthen external validity of the hypotheses H2 and H3, formal hypothesis testing in a controlled experiment could be pursued.

(2) **Utility of BluGen to SAs**. The perceived utility of BluGen to working SAs could be assessed using survey methods.

(3) **Assess Utility of Mitigation Recommendations**. Experimental tests of the degree to which implementations of the mitigation recommendations from BluGen hold up against anticipated threat actors could be evaluated.

(4) **Explore Empirical Validation of BluGen RefCat**. One could evaluate the process of empirical validation of RefCat contents using actual cyber incident data.

(5) **Willingness to Review and Contribute**. A study to examine the extent to which the broader cyber community is willing to reuse, contribute to, and peer review BluGen RefCat content could be undertaken.

(6) **Use of BluGen for other Threat Types**. The expansion of BluGen to other threat types besides cyber (e.g., kinetic threats, electromagnetic threats) could be examined. At issue would be how well BluGen analytics and BluGen's capability-based representation of threats and mitigations work.

(7) **Real-Time BluGen**. An examination of the degree to which BluGen could be extended to do "real-time" risk assessment could be undertaken. Such a tool could be driven by data from live update feeds of threat data and system configuration data.

(8) **Tradespace Analysis.** Tradespace analysis of possible capability-based mitigation architectures is a rich area for possible future investigation. In this context, one could build a recommendation engine that selects mitigations not just on the basis of the perceived effectiveness of individual defensive solutions, but on the effectiveness of overall mitigation architectures composed of those solutions, taking into consideration various SA-weighted measures of cost and benefit.

REFERENCES

Alavi, M., & Leidner, D. (2001). Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues. *MIS Quarterly*, *25*(1), 107–136. Retrieved from http://www.jstor.org/stable/3250961

Aven, T. (2016). Risk assessment and risk management: Review of recent advances on their foundation. *European Journal of Operational Research*, *253*(1), 1–13. http://doi.org/http://dx.doi.org/10.1016/j.ejor.2015.12.023

Bateman, A. (2017). In outer space, the US is vulnerable to China and Russia. *The Hill*. Retrieved from http://thehill.com/blogs/pundits-blog/defense/342992-in-outer-space-the-us-is-vulnerable-to-china-and-russia

Becerra-Fernandez, I., & Sabherwal, R. (2010). *Knowledge Management: Systems and Processes*. M.E.Sharpe.

Bolger, F., & Wright, G. (1994). Assessing the quality of expert judgment: Issues and analysis. *Decision Support Systems*, *11*(1), 1–24. http://doi.org/https://doi.org/10.1016/0167-9236(94)90061-2

Caralli, R., Stevens, J., Young, L., & Wilson, W. (2007). Introducing OCTAVE Allegro: Improving the Information Security Risk Assessment Process. Retrieved from http://www.cert.org/octave

Carlson, C., Hutton, A., & Gilliam, A. (2010). *FAIR – ISO/IEC 27005 Cookbook - Technical Guide*. Retrieved from http://www.businessofsecurity.com/docs/FAIR - ISO_IEC_27005 Cookbook.pdf

CNSS Instruction No. 1253 - Security Categorization and Control Selection for National Security Systems, Version 2. (2012). Retrieved from http://www.sandia.gov/FSO/PDF/flowdown/Final_CNSSI_1253.pdf

Committee on National Security Systems. (2010). CNSS 4009 - National Information Assurance (IA) Glossary. Retrieved from http://www.ncsc.gov/publications/policy/docs/CNSSI_4009.pdf

Congress. (2016). H.R.1735 - National Defense Authorization Act. Retrieved from https://www.congress.gov/bill/114th-congress/house-bill/1735

Cybersecurity Ventures. (2016). Cybersecurity Market Report. Retrieved from http://cybersecurityventures.com/cybersecurity-market-report

Dewri, R., Poolsappasit, N., Ray, I., & Whitley, D. (2007). Optimal Security Hardening Using Multi-objective Optimization on Attack Tree Models of Networks. In *ACMCCS*. Retrieved from http://www.cs.colostate.edu/~genitor/2007/2007ACMCCS.pdf

Dinsmore, P. (2016). NIPRNet/SIPRNet Cyber Security Architecture Review. Retrieved from http://www.disa.mil/~/media/Files/DISA/News/Conference/2016/AFCEA-Symposium/3-Dinsmore_NSCSAR.pdf

DoD. (2015). *DoD CIO/AT&L Capability-based Threat Model*.

Goodall, J. R., D'Amico, A., & Kopylec, J. K. (2009). Camus: Automatically mapping Cyber Assets to Missions and Users. *Military Communications Conference, 2009. MILCOM 2009. IEEE*. http://doi.org/10.1109/MILCOM.2009.5380096

Gosler, J., & Von Thaer, L. (2013). *Resilient Military Systems and the Advanced Cyber Threat*. Retrieved from http://www.acq.osd.mil/dsb/reports/ResilientMilitarySystems.CyberThreat.pdf

Hallberg, J., Bengtsson, J., Hallberg, N., Karlzén, H., & Sommestad, T. (2017). The Significance of Information Security Risk Assessments Exploring the Consensus of Raters' Perceptions of Probability and Severity. In *International Conference on Security and Management* (pp. 131–137).

Heater, B. (2017). Comparing Alexa, Google Assistant, Cortana and Siri smart speakers. Retrieved from https://techcrunch.com/2017/10/08/comparing-alexa-google-assistant-cortana-and-siri-smart-speakers

Hevner, B. A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly*, *28*(1), 75–105.

Holm, H., Sommestad, T., Ekstedt, M., & Honeth, N. (2014). Indicators of expert judgement and their significance: An empirical investigation in the area of cyber security. *Expert Systems*, *31*(4), 299–318.

House Government Reform Committee. (2002). H.R.3844 - Federal Information Security Management Act of 2002. Retrieved from https://www.congress.gov/bill/107th-congress/house-bill/3844

Hubbard, D. W., & Seiersen, R. (2016). *How to Measure Anything in Cybersecurity Risk*. Wiley.

INFOSEC Institute. (2013). Quantitative Risk Analysis. Retrieved from http://resources.infosecinstitute.com/quantitative-risk-analysis/#gref

InsurTech. (2017). Cyber Crime Risk Management. Retrieved from https://justinsurtech.com/cyber-crime-risk-management

International Standards Organization. (2009). ISO 31000 - Risk management.

Investopedia Staff. (n.d.). Financial Concepts: The Risk/Return Tradeoff. Retrieved from http://www.investopedia.com/university/concepts/concepts1.asp

ISO/IEC 27001:2013 - Information technology, Security techniques, Information security Management systems, Requirements. (2013). Retrieved from http://www.iso.org/iso/catalogue_detail?csnumber=54534

Kaplan, S. (1997). *The Words of Risk Analysis*. Retrieved from http://onlinelibrary.wiley.com/doi/10.1111/j.1539-6924.1997.tb00881.x/abstract

Kaplan, S., & Garrick, B. J. (1981). On the Quantitative Definition of Risk. *Society for Risk Analysis*, *1*(1), 17. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.323.1418&rep=rep1&type=pdf

Llanso, T., & Engebretson, P. (2016). A Unified Model for System Security Engineering. In *Hawaii International Conference on System Sciences* (p. 8).

Llanso, T., & Klatt, E. (2014). CyMRisk: An approach for computing mission risk due to cyber attacks. In *IEEE International Systems Conference. Ottawa* (pp. 1–7). http://doi.org/10.1109/SysCon.2014.6819227

Llanso, T., & McNeil, M. (2018). Estimating Software Vulnerability Counts in the Context of Cyber Risk Assessments. In *Hawaii International Conference on System Sciences* (p. 7).

Llanso, T., McNeil, M., Pearson, D., & Moore, G. (2017). An Analytic Framework for Mission-Cyber Risk Assessment and Mitigation Recommendation. In *Hawaii International Conference on System Sciences* (p. 10). Retrieved from http://www.hicss.org

Lukyanenko, R., Evermann, J., & Parsons, J. (2014). Instantiation Validity in IS Design Research. In M. C. Tremblay, D. VanderMeer, M. Rothenberger, A. Gupta, & V. Yoon (Eds.), *Advancing the Impact of Design Science: Moving from Theory to Practice: 9th International Conference, DESRIST 2014, Miami, FL, USA, May 22-24, 2014. Proceedings* (pp. 321–328). Cham: Springer International Publishing. http://doi.org/10.1007/978-3-319-06701-8_22

McLeod, S. (2008). Likert Scale. Retrieved from http://www.simplypsychology.org/likert-scale.html

McNeil, M., Llanso, T., & Pearson, D. (2018). Application of Capability-Based Cyber Risk Assessment Methodology to a Space System. In *Hot Topics in the Science of Security Symposium*.

Merriam-Webster. (n.d.). “Cyber” Definition. Retrieved from https://www.merriam-webster.com/dictionary/cyber

Merriam-Webster Online Dictionary. (n.d.). Retrieved from https://www.merriam-webster.com/

Mitre. (n.d.). Common Attack Pattern Enumeration and Classification. Retrieved from https://capec.mitre.org/index.html

Musman, S., Tanner, M., Temin, A., Elsaesser, E., & Loren, L. (2011). Computing the Impact of Cyber Attacks on Complex Missions. *2011 IEEE International Systems Conference*, 46–51. http://doi.org/10.1109/SYSCON.2011.5929055

National Institute of Standards and Technology. (2010). *Guide for Applying the Risk Management Framework to Federal Information Systems*. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf

National Institute of Standards and Technology. (2012). *National Institute of Standards and Technology 800-30: Guide for Conducting Risk Assessments*. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-30-rev1/sp800_30_r1.pdf

National Institute of Standards and Technology (NIST). (2004). *FIPS Pub 199 - Standards for Security Categorization of Federal Information and Information Systems*. Retrieved from http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf

National Institute of Standards and Technology Special Publication 800-53 Revision 4. (2013). Retrieved from http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf

Newhouse, W., Keith, S., Scribner, B., & Witte, G. (2017). *NIST Special Publication 800-181 - National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework*. Retrieved from http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-181.pdf

NIST. (2010). NIST Special Publication 800-37, Revision 1: Guide for Applying the Risk Management Framework to Federal Information Systems. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf

NVD. (n.d.). National Vulnerability Database. Retrieved from http://nvd.nist.gov

O'Dell, C., & Hubert, C. (2011). *New Edge in Knowledge: How Knowledge Management Is Changing the Way We Do Business*. Wiley.

Object Management Group. (1999). Unified Modeling Language (UML). Retrieved from http://www.uml.org

Orfei, S., Leach, T., King, J., Mauro, L., & Fitzsimmons, J. (2006). Payment Card Industry Security. Retrieved from https://www.pcisecuritystandards.org/pci_security/

Peacos, P. (2016). Bias: The Hidden Danger to Your Risk Assessment. Retrieved from http://www.americanpharmaceuticalreview.com/Featured-Articles/184365-Bias-The-Hidden-Danger-to-Your-Risk-Assessment/

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. *J. Manage. Inf. Syst.*, *24*(3), 45–77. http://doi.org/10.2753/MIS0742-1222240302

Rausand, M. (2011). *Risk Assessment: Theory, Methods, and Applications*. John Wiley &

Sons, Inc.

Stine, K., Kissel, R., Barker, W., Fahlsing, J., & Gulick, J. (2008). Volume I: Guide for Mapping Types of Information and Information Systems to Security Categories. Retrieved from http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-60v1r1.pdf

Trochim, W., & Donnelly, J. (2008). *The Research Methods Knowledge Base*. Atomic Dog.

US Department of Defense. (2014). Department of Defense Instruction Number 8510.01 - Risk Management Framework (RMF) for DoD Information Technology (IT). Retrieved from http://www.dtic.mil/whs/directives/corres/pdf/851001_2014.pdf

Vaishnavi, V., & Kuechler, B. (2011). Design Science Research in Information Systems. Retrieved from http://desrist.org/desrist/content/design-science-research-in-information-systems.pdf

Various. (n.d.). Society for Risk Analysis. Retrieved from http://www.sra.org

Vigo, R., Nielson, F., & Nielson, H. R. (2014). Automated Generation of Attack Trees. In *Proceedings of the 2014 IEEE 27th Computer Security Foundations Symposium* (pp. 337–350). Washington, DC, USA: IEEE Computer Society. http://doi.org/10.1109/CSF.2014.31

Wynn, J., Whitmore, J., Upton, G., & Spriggs, L. (2011). *Threat Assessment & Remediation Analysis (TARA)*. Retrieved from https://www.mitre.org/sites/default/files/pdf/11_4982.pdf

Xia, F., Yang, L., Wang, L., & Vinel, A. (2012). Internet of Things. *International Journal of Communications Systems*, *25*, 1101–1102.

Yevseyeva, I., Basto-Fernandes, V., Emmerich, M., & van Moorsel, A. (2015). Selecting Optimal Subset of Security Controls. *Procedia Computer Science*, *64*, 1035–1042. http://doi.org/https://doi.org/10.1016/j.procs.2015.08.625

APPENDICES

Appendix A - Additional Information on BluGen

Appendix A provides additional information on BluGen, broken into two sections: (1) setting up and running BluGen and (2) BluGen data capture. We do not summarize BluGen itself, as that was done in the earlier section called Artifact Design and in the 2017 HICSS paper (Llanso et al., 2017).

Setting Up and Running the BluGen Software

BluGen software is managed in the SVN repository, which should already be installed. To check out the software, follow these steps:

1. Create an empty folder. Below, I called it BluGen-1.0Demo
2. Change directories to the folder
3. Right-click mouse and choose "SVN Checkout…"

4. The dialog box below appears. Enter the appropriate URL and then click OK. The checkout process will commence.

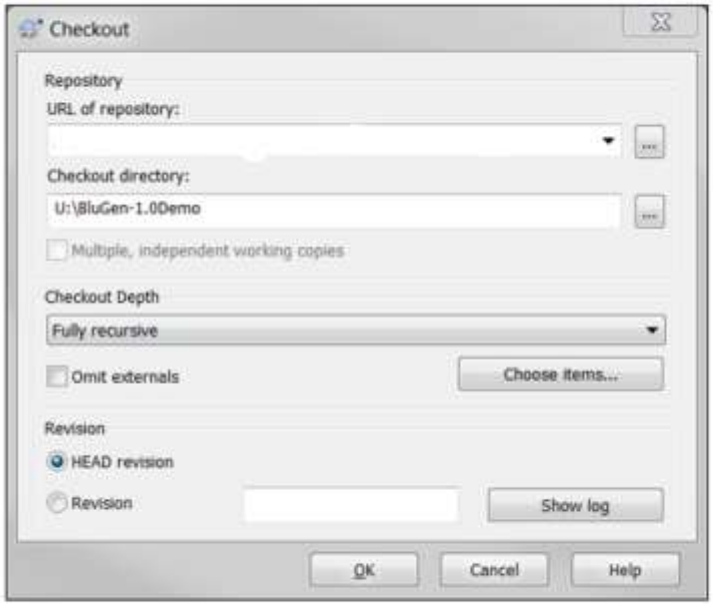

5. The checkout process takes several minutes. When the process is complete, the following window contents will appear.

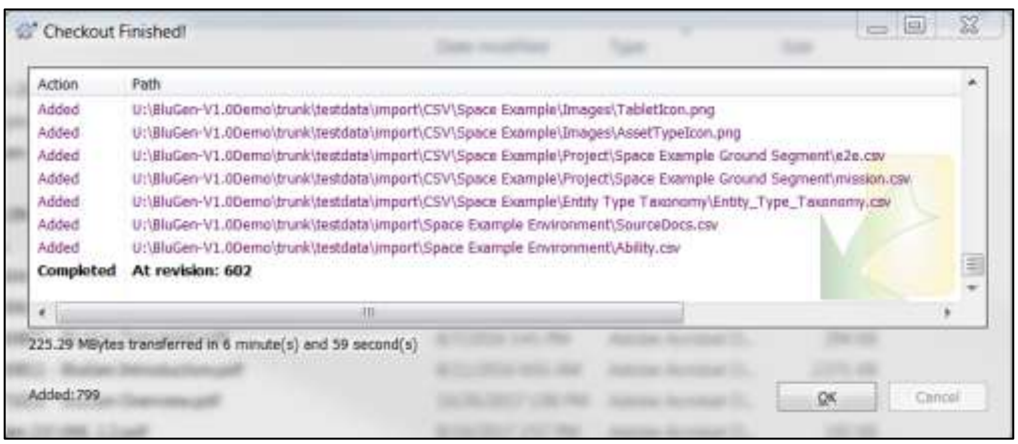

The folder appears as follows once the checkout is complete.

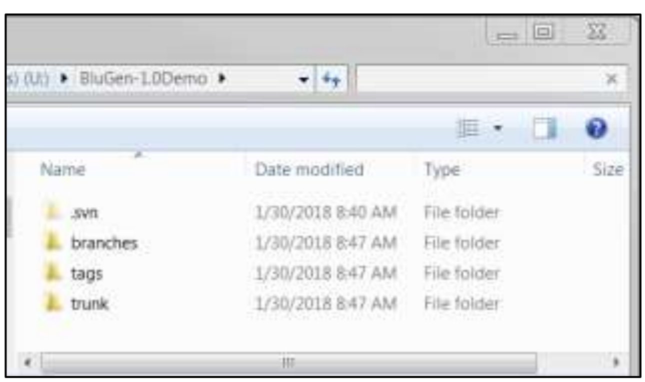

6. Start the NetBeans IDE (version 8.2 used below), choose Open Project… from the File menu, which results in the dialog below. Then choose “trunk” in the file list.

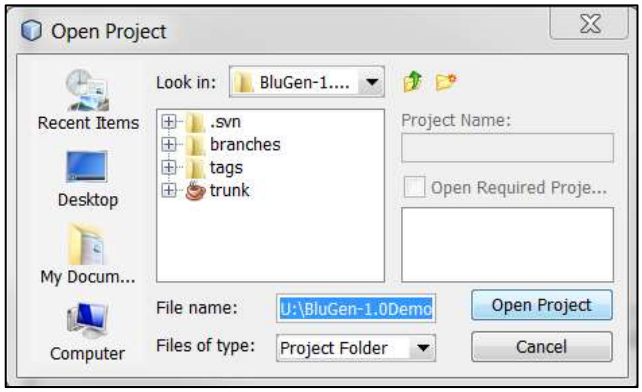

For the sake of brevity, we do not show the installation process for installing the BluGen software nor do we show the importation process for Omega descriptive data[10].

To start the BluGen tool, the user double-clicks the mouse on the BluGen icon on the desktop (Figure 19).

Figure 19: Desktop with BluGen icon

The tool starts up and presents the user with a list of projects (Figure 20). A project is a description of a target system to be analyzed. Omega has already been loaded into a project called “Space Example”.

[10] To import the data, the user prepares a multi-tab spreadsheet populated with descriptive data for Omega. The user then executes a command in BluGen to load this data into a newly created BluGen project.

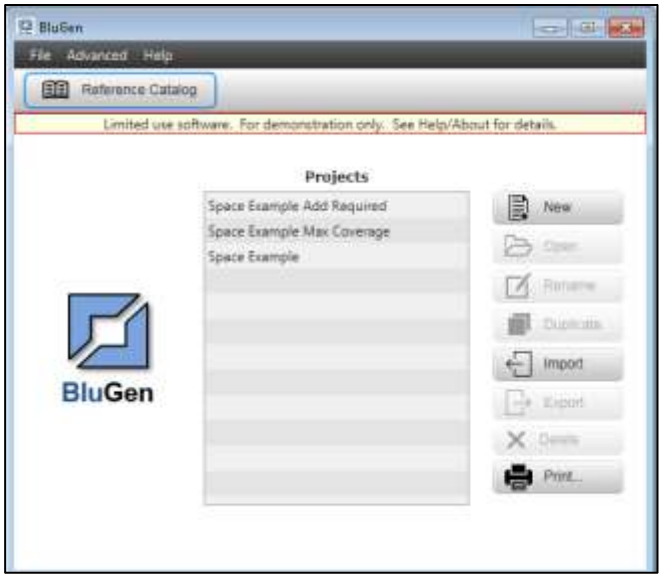

Figure 20: BluGen Projects

To view details about the risk analysis of Omega, the user selects the "Space Example" project with the mouse and clicks the "Open" button. The corresponding project window opens (Figure 21). Note the multi-tab interface for the project description. The main tab, shown below, captures the project name, description, threat model to use, tier of threat actor to consider, and risk tolerance values.

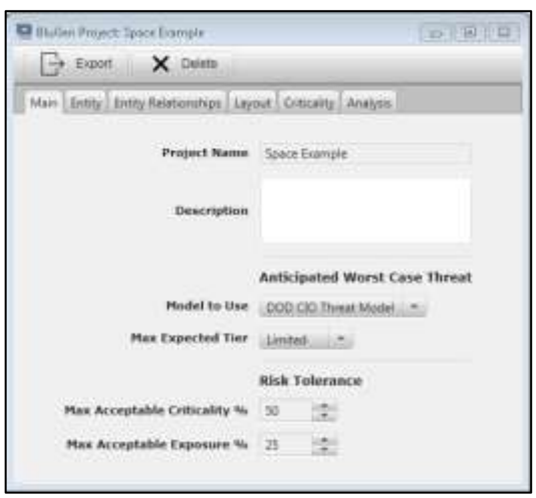

Figure 21: Project Windows – Main tab

Figure 22 shows a view of the entity tab for Omega. Entities include missions, assets, and data types. The window shows only a subset of the entities in Omega.

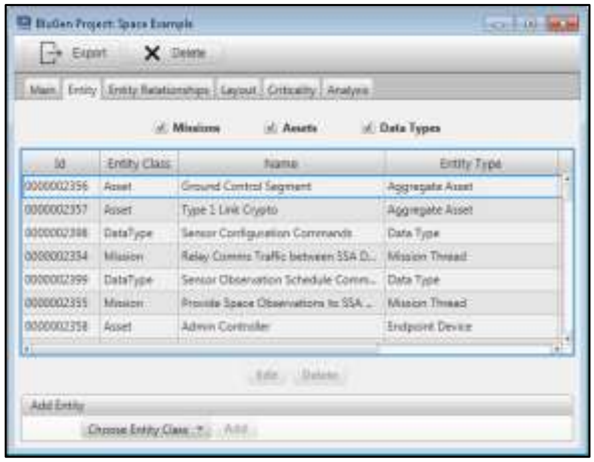

Figure 22: Project Windows – Entity Tab

Figure 23 shows a view of the entity relationships tab for Omega. For example, one of the entity relationship types is "InheritsCapabilitiesFrom," which indicates that an asset inherits the defensive mitigations from another asset.

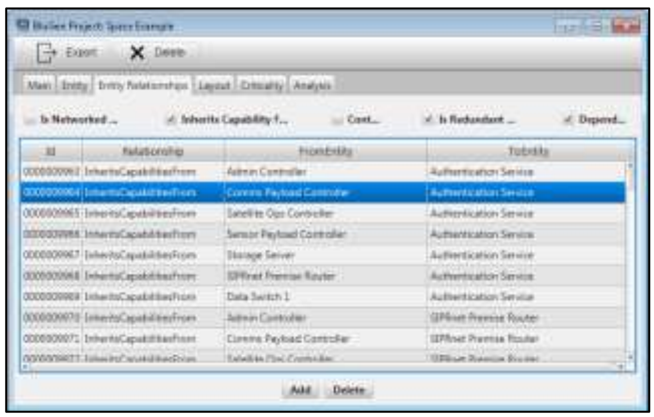

Figure 23: Project Windows – Entity Relationships Tab

Figure 24 shows the mission criticality scores for Omega. The user provides this data as input to BluGen. Each row of data in the table shows mission impact scores for breach of confidentiality, integrity, and availability for each viable combination of (Mission, Asset, and Data).

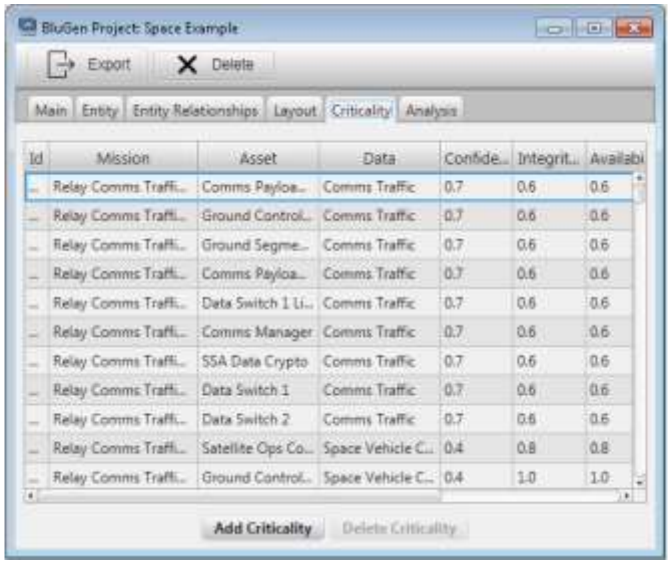

Figure 24: Project Windows – Criticality Tab

Figure 25 shows the analysis tab for Omega. The analyst clicks on the various buttons to run BluGen analytics. The buttons used for this analysis are the buttons to generate a risk plot and to generate a mitigations report.

Figure 25: Project Windows – Analysis Tab

Figure 26 show a risk plot generated for Omega. Each data point in the scatterplot represents an asset in the Project model provided to BluGen.

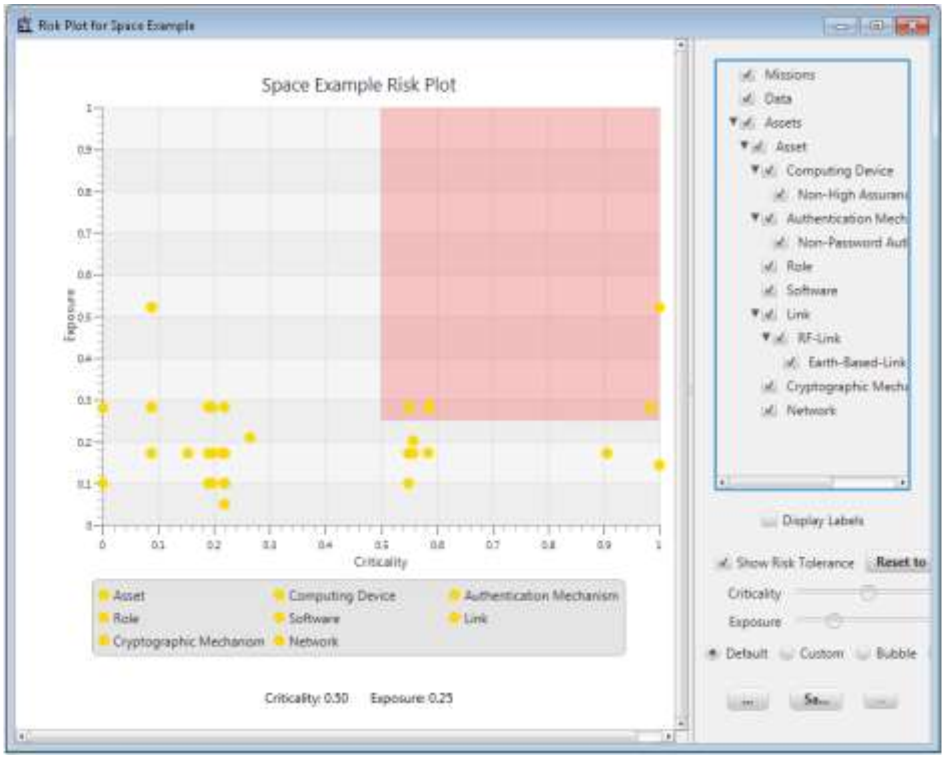

Figure 26: Risk Plot Generated for Omega

Figure 27 shows a portion of the mitigation report BluGen generates for Omega. Each row in the table represents an asset. Mitigation possibilities for the asset are shown on the right-hand side of the report.

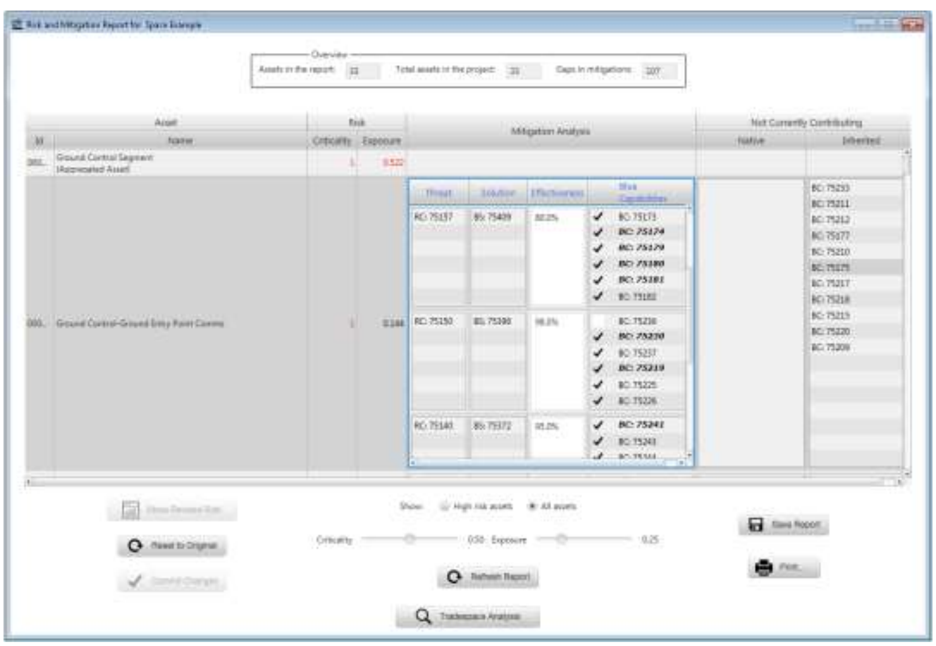

Figure 27: Mitigation Report Generated for Omega

Omega Data Capture

This section discusses dissertation data capture. In this context, by data, we mean processing data generated by the BluGen tool needed for evaluating hypothesis H3. We first show the custom code that we wrote to extract the data, and then we show the data itself.

Custom Source Code for Dissertation Hypothesis H3 (BluGen)

For the dissertation, I wrote custom source code to capture data processed by the BluGen exposure analytic as it worked its way through computing exposure for assets in Omega. This code, which is not part of the main BluGen source code base, is shown in below.

```
package jhuapl.edu.blugen.ui;

import java.util.ArrayList;
import java.util.HashMap;
import jhuapl.edu.blugen.EntityTypeTaxonomyManager;
import jhuapl.edu.blugen.ReferenceCatalogManager;
import jhuapl.edu.blugen.model.EntityType;
import jhuapl.edu.blugen.model.EntityTypeTaxonomy;
import jhuapl.edu.blugen.model.refcat.Ability;
import jhuapl.edu.blugen.model.refcat.Ability2Ability;
import jhuapl.edu.blugen.model.refcat.ReferenceCatalog;

/**
 * This code was written by Thomas H. Llanso in support of his dissertation.
 *
 * @author Thomas H. Llanso
 */
public class Dissertation {

    /**
     * This method traverses the reference catalog for each asset type found in the
     * Space Example project, showing coverage and collecting descriptive
     * statistics along the way.
     */
    public void execute() {
        EntityTypeTaxonomyManager ettm = EntityTypeTaxonomyManager.getInstance();
        EntityTypeTaxonomy att = ettm.getEntityTypeTaxonomy();
        ReferenceCatalogManager rcm = ReferenceCatalogManager.getInstance();
        ReferenceCatalog rc = rcm.getReferenceCatalog();

        System.out.println("******* BluGen Dissertation Output *******");

        int assetTypeCount = 0;
        ArrayList<Ability> rcList = new ArrayList<>();
        ArrayList<Ability> bsList = new ArrayList<>();
        ArrayList<Ability> bcList = new ArrayList<>();
        int[] mappings = new int[3];

        for (EntityType et : att.getEntityTypes().values())
            if (dissertation_assetTypeWasUsed(et)) {
                dissertation_ShowCoverageForEntityType(rc, et, rcList, bsList,
                        bcList, mappings);
                assetTypeCount++;
```

```
            }
        System.out.println("\n         ---> STATISTICS <---");
        System.out.println("           Asset Types (AT) count: "+assetTypeCount);
        System.out.println("   Offensive Capability (OC) Count: "+rcList.size());
        System.out.println("     Defensive Solution (DS) Count: "+bsList.size());
        System.out.println("   Defensive Capability (DC) Count: "+bcList.size());
        System.out.println("");
        System.out.println("          OC --> AT Mapping Count: "+mappings[0]);
        System.out.println("          BS --> RC Mapping Count: "+mappings[1]);
        System.out.println("          BC --> BS Mapping Count: "+mappings[2]);
    }

/**
 * THis method returns TRUE if a given asset type was used in the Space Example.
 *
 * @param et Asset type to lookup
 * @return TRUE if present, FALSE if not.
 */
boolean dissertation_assetTypeWasUsed(EntityType et) {
    // Entity types in Space Example
    String[] aList = {
        "Aggregate Asset",
        "Authentication Mechanism",
        "Computing Device",
        "Endpoint Cryptographic Mechanism",
        "Endpoint Device",
        "General User",
        "Key Management Mechanism",
        "Network Device",
        "Non-IT Roles",
        "Physical Space",
        "Security Admin Roles",
        "System Admin Roles",
        "Wired-Link"
    };
    boolean found = false;
    for (String name : aList) {
        if (name.equalsIgnoreCase(et.getName())) {
            found = true;
            break;
        }
    }
    return found;
}

    /**
     * Show the coverage for a given asset type.
     *
     * @param rc Reference Catalog to use
     * @param entityType Asset type to show coverage for
     * @param rcList Accumulating list of offensive capabilities
     * @param bsList Accumulating list of defensive solutions
     * @param bcList Accumulating list of defensive capabilities
     * @param mappings Accumulating list of mappings between entities
     */
    void dissertation_ShowCoverageForEntityType(
            ReferenceCatalog rc,
            EntityType entityType,
            ArrayList<Ability> rcList,
            ArrayList<Ability> bsList,
            ArrayList<Ability> bcList,
            int[] mappings) {
```

```
        System.out.println("\nASSET-TYPE: "+entityType.getName());

        // Show red capabilities and corresponding blue solutions and component blue capabilities
        for (Ability redAbility : rc.getRedAbilitiesThatThreatenEntityType(entityType, null, true)) {
            dissertation_addAbility(rcList, redAbility);
            mappings[0]++;

            System.out.println("      OC: "+dissertation_trim(redAbility.getName()));
            HashMap<Ability, Double> map = rc.getBlueAbilitiesThatCounterRedAbility(redAbility, null);
            for (Ability bs : map.keySet()) {
                System.out.println("            DS: "+
                    dissertation_trim(bs.getName().substring(4))); //+" ("+bs.getAbilityCategory()+")");
                mappings[1]++;
                dissertation_addAbility(bsList, bs);
                ArrayList<Ability2Ability> list = rc.getComposedOf(bs);
                for (Ability2Ability a2a : list) {
                    Ability bc = a2a.getAbility2();
                    mappings[2]++;
                    dissertation_addAbility(bcList, bc);
                    System.out.println("                  DC: "+
                        dissertation_trim(bc.getName().substring(4))); //+" ("+bc.getAbilityCategory()+")");
                }
            }
        }
    }

    /**
     * Trim output string to no longer than 100 characters
     *
     * @param s String to trim
     * @return s trimmed string
     */
    String dissertation_trim(String s) {
        final int m = 115;
        int len = s.length() > m ? m : s.length();
        String k = s.substring(0, len);
        if (k.length() == m) k += "...";
        return k;

    }

    /**
     * Add an ability to the list as long as it is not already on the list.
     *
     * @param list list to receive the ability
     * @param a  ability
     */
    void dissertation_addAbility(ArrayList<Ability> list, Ability a) {
        boolean found = false;
        for (Ability i : list)
            if (i.getName().equalsIgnoreCase(a.getName())) {
                found = true;
                break;
            }
        if (!found)
            list.add(a);
    }
}
```

To summarize the code above, for each asset type that appears in Omega, the code shows the offensive capabilities (OC) mapped to the asset types, and, for each threat, the defensive solutions (DS) that mitigate the threat, and the defensive capabilities (DC) that compose those solutions. In addition, the code computes summary statistics at the very end.

BluGen Output to Show Coverage

A sampling of the output resulting from a run of the Java code for Omega is shown for below. The descriptive statistics that normally appear at the end of the multi-page output is instead show in Figure 28 for convenience.

```
    ---> STATISTICS <---
      Asset Types (AT) count: 13
Offensive Capability (OC) Count: 48
 Defensive Solution (DS) Count: 86
Defensive Capability (DC) Count: 47

     OC --> AT Mapping Count: 129
     BS --> RefCat Mapping Count: 303
     BC --> BS Mapping Count: 383
```

Figure 28: BluGen Descriptive Statics for Omega Coverage

```
******* BluGen Dissertation Output *******

ASSET-TYPE: General User
  OC: Effectively uses highly-sophisticated techniques in social settings for elicitation
    DS: Solution to Highly Sophisticated Social Engineering Full GR
      DC: Mitigate Highly Sophisticated Social Mining for Elicitation G
    DS: Solution to Highly Sophisticated Social Mining for Elicitation G
      DC: Mitigate Highly Sophisticated Social Mining for Elicitation R
    DS: Solution to Highly Sophisticated Social Mining for Elicitation R
      DC: Mitigate Highly Sophisticated Social Mining for Elicitation GR
  OC: Effectively uses highly sophisticated social engineering attacks
    DS: Solution to Highly Sophisticated Social Engineering R
      DC: Mitigate Moderately Sophisticated Social Mining for Elicitation G
  OC: Effectively uses moderately-sophisticated techniques in social settings for elicitation
    DS: Solution to Moderately Sophisticated Social Mining for Elicitation Full GR
      DC: Mitigate Moderately Sophisticated Social Mining for Elicitation G
    DS: Solution to Moderately Sophisticated Social Mining for Elicitation R
      DC: Mitigate Moderately Sophisticated Social Mining for Elicitation GR
    DS: Solution to Moderately Sophisticated Social Mining for Elicitation G
      DC: Mitigate Moderately Sophisticated Social Mining for Elicitation R
  OC: Effectively uses moderately sophisticated techniques to recruit persons for espionage and sabot...
    DS: Mitigate Moderately Sophisticated Recruitment Techniques
      DC: Mitigate Moderately Sophisticated Recruitment Techniques
  OC: Effectively uses highly sophisticated techniques to recruit persons for espionage and sabotage
    DS: Mitigate Highly Sophisticated Recruitment Techniques
      DC: Mitigate Highly Sophisticated Recruitment Techniques
  OC: Effectively uses rudimentary social engineering attacks
    DS: Solution to Rudimentary Social Engineering R
      DC: Mitigate Rudimentary Social Engineering R
    DS: Solution to Rudimentary Social Engineering Full GR
      DC: Mitigate Rudimentary Social Engineering GR
```

DS: Solution to Rudimentary Social Engineering G
DC: Mitigate Rudimentary Social Engineering G
OC: Effectively uses moderately sophisticated social engineering attacks
DS: Solution to Moderately Sophisticated Social Engineering G
DC: Mitigate Moderately Sophisticated Social Engineering G
DS: Solution to Moderately Sophisticated Social Engineering R
DC: Mitigate Moderately Sophisticated Social Engineering R
DS: Solution to Moderately Sophisticated Social Engineering Full GR
DC: Mitigate Moderately Sophisticated Social Engineering GR

ASSET-TYPE: Endpoint Cryptographic Mechanism
OC: Defeat Commercial Crypto using cryptanalysis
DS: Mitigate ability to defeat Commercial Crypto via military Grade Encryption
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
OC: Defeat a weak commercial cryptographic mechanism in a computing device
DS: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
DC: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
OC: Defeats Strong Commercial Crypto by obtaining key material
DS: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
DC: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms that are unprotected
DS: Can mitigate attacks on data on RF Links that are unprotected using cryptography
DC: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
DS: Can mitigate attacks on data on RF Links that are unprotected using cryptography and physical protections
DC: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...
DC: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
DS: Can mitigate attacks on data on RF Links that are unprotected using physical access controls
DC: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with simple protections
DS: Can mitigate attacks on data on RF Links with simple protections using physical access controls
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 2, or less) with ...
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with highly-sophisticated protec...
DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using cryptography
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material
DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using cryptography and physical prot...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection, classifi...
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)
DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using physical access controls
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with extra highly-sophisticated ...
DS: Can mitigate attacks on data on RF Links with extra highly-sophisticated protections using physical access controls...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DS: Can mitigate attacks on data on RF Links with extra highly-sophisticated protections using cryptography
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)
OC: Defeats Military Grade Crypto by obtaining key material
DS: Mitigate ability to defeat Military Grade Crypto by obtaining key material
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with rudimentary protections
DS: Can mitigate attacks on data on RF Links with rudimentary protections using physical access controls
DC: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
DS: Can mitigate attacks on data on RF Links with rudimentary protections using cryptography
DC: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
DS: Can mitigate attacks on data on RF Links with simple protections using cryptography and physical protections
DC: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
DC: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
OC: Defeats Military Grade Crypto by obtaining key material (faster than T5)
DS: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with moderately-sophisticated pr...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography reducing vuln...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption

DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls a...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls b...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls c...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography using stronge...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with rudimentary protections using cryptography and physical protections a...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis

ASSET-TYPE: Network Device
OC: Use compromised humans to attack data on devices and links through highly-sophisticated social engineering/elicitat...
DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated social engineer...
DC: Mitigate Highly Sophisticated Social Mining for Elicitation G
DC: Mitigate Highly Sophisticated Social Engineering GR
DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated social engineer...
DC: Mitigate Highly Sophisticated Social Mining for Elicitation R
DC: Mitigate Highly Sophisticated Social Engineering GR
DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated social engineer...
DC: Mitigate Highly Sophisticated Social Mining for Elicitation GR
DC: Mitigate Highly Sophisticated Social Engineering GR
OC: Exploit Known and Unknown Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
DS: BS: Detect and Respond to exploitation of Known Weak Configurations Settings (CCEs) in Software (OS, firmware, Appl...
DC: BG: Protect against Known Weak Configurations Settings (CCEs)
OC: Use compromised humans to attack data on devices and links through basic recruitment through moderately-sophisticat...
DS: Mitigate use of compromised humans to attack data on devices and links through basic recruitment through moderately...
DC: Mitigate Moderately Sophisticated Recruitment Techniques
OC: Use compromised humans to attack data on devices and links through rudimentary social engineering
DS: Mitigate use of compromised humans to attack data on devices and links through rudimentary social engineering G
DC: Mitigate Rudimentary Social Engineering G
DS: Mitigate use of compromised humans to attack data on devices and links through rudimentary social engineering R
DC: Mitigate Rudimentary Social Engineering R
DS: Mitigate use of compromised humans to attack data on devices and links through rudimentary social engineering GR
DC: Mitigate Rudimentary Social Engineering GR
OC: Exploit Known and Unknown Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
DS: BS: Protect against Known Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
DC: BG: Protect against Known Weak Configurations Settings (CCEs)
OC: Use compromised humans to attack data on devices and links through moderately-sophisticated social engineering/elic...
DS: Mitigate use of compromised humans to attack data on devices and links through moderately-sophisticated social engi...
DC: Mitigate Moderately Sophisticated Social Engineering GR
DC: Mitigate Moderately Sophisticated Social Mining for Elicitation GR
DS: Mitigate use of compromised humans to attack data on devices and links through moderately-sophisticated social engi...
DC: Mitigate Moderately Sophisticated Social Engineering G
DC: Mitigate Moderately Sophisticated Social Mining for Elicitation G
DS: Mitigate use of compromised humans to attack data on devices and links through moderately-sophisticated social engi...
DC: Mitigate Moderately Sophisticated Social Engineering R
DC: Mitigate Moderately Sophisticated Social Mining for Elicitation R
OC: Exploit Known and Unknown Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
DS: BS: Limit damage from Known Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor)...
DC: BG: Protect against Known Weak Configurations Settings (CCEs)
OC: Can develop and deliver high-stealth SW implants for SW of network appliances and embedded systems
DS: Mitigate SW Injection: SW Hash-Based WL TT4

DC: Mitigate Malicious and Unauthorized Code emphasis Hash based Whitelisting to block execution
DS: Mitigate SW Injection: SW Black Listing TT4
DC: Mitigate Malicious and Unauthorized Code emphasis Black Listing
DS: Mitigate SW Injection: SW Location WL + Hash-based Removal TT4
DC: Mitigate Malicious and Unauthorized Code emphasis Locational WL with Hash based Removal of Malicious Code
DS: Mitigate SW Injection: SW Locational WL TT4
DC: Mitigate Malicious and Unauthorized Code emphasis Location Whitelisting to block execution
OC: Exploit Known Vulnerabilities (CVEs and CWEs) in Software (OS, firmware, Application, Hypervisor) of computers, sma...
DS: Mitigate Exploitation of known Vulnerabilities CVEs and CWEs) in Software (OS, firmware, Application, Hypervisor) o...
DC: Mitigate Exploitation of known Vulnerabilities CVEs and CWEs) in Software (OS, firmware, Application, Hypervisor) o...
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms that are unprotected
DS: Can mitigate attacks on data on RF Links that are unprotected using cryptography
DC: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
DS: Can mitigate attacks on data on RF Links that are unprotected using cryptography and physical protections
DC: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...
DC: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
DS: Can mitigate attacks on data on RF Links that are unprotected using physical access controls
DC: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with simple protections
DS: Can mitigate attacks on data on RF Links with simple protections using physical access controls
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 2, or less) with ...
OC: Use compromised humans to attack data on devices and links through highly-sophisticated recruitment for espionage/...
DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated recruitment fo...
DC: Mitigate Highly Sophisticated Recruitment Techniques
OC: Inject Hardware
DS: Mitigate hardware injection
DC: Mitigate Hardware Injection
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with highly-sophisticated protec...
DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using cryptography
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material
DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using cryptography and physical prot...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection, classifi...
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)
DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using physical access controls
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with extra highly-sophisticated ...
DS: Can mitigate attacks on data on RF Links with extra highly-sophisticated protections using physical access controls...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DS: Can mitigate attacks on data on RF Links with extra highly-sophisticated protections using cryptography
DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with rudimentary protections
DS: Can mitigate attacks on data on RF Links with rudimentary protections using physical access controls
DC: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
DS: Can mitigate attacks on data on RF Links with rudimentary protections using cryptography
DC: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
DS: Can mitigate attacks on data on RF Links with simple protections using cryptography and physical protections
DC: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
DC: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
OC: Find and Exploit Unknown Vulnerabilities in OS, firmware or application software on computing devices
DS: Mitigate Exploitation of unknown Vulnerabilities in OS, firmware or application software on computing devices
DC: Mitigate Exploitation of unknown Vulnerabilities in OS, firmware or application software on computing devices
DC: Mitigate Exploitation of unknown Vulnerabilities in hypervisor software on computing devices
OC: Exploit Hardware Vulnerabilities
DS: Mitigate hardware vulnerability
DC: Mitigate Vulnerable Hardware
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with moderately-sophisticated pr...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography reducing vuln...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls a...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls b...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...

DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls c...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography using stronge...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with rudimentary protections using cryptography and physical protections a...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis

ASSET-TYPE: System Admin Roles
OC: Effectively uses highly-sophisticated techniques in social settings for elicitation
DS: Solution to Highly Sophisticated Social Engineering Full GR
DC: Mitigate Highly Sophisticated Social Mining for Elicitation G
DS: Solution to Highly Sophisticated Social Mining for Elicitation G
DC: Mitigate Highly Sophisticated Social Mining for Elicitation R
DS: Solution to Highly Sophisticated Social Mining for Elicitation R
DC: Mitigate Highly Sophisticated Social Mining for Elicitation GR
OC: Effectively uses highly sophisticated social engineering attacks
DS: Solution to Highly Sophisticated Social Engineering R
DC: Mitigate Moderately Sophisticated Social Mining for Elicitation G
OC: Effectively uses moderately-sophisticated techniques in social settings for elicitation
DS: Solution to Moderately Sophisticated Social Mining for Elicitation Full GR
DC: Mitigate Moderately Sophisticated Social Mining for Elicitation G
DS: Solution to Moderately Sophisticated Social Mining for Elicitation R
DC: Mitigate Moderately Sophisticated Social Mining for Elicitation GR
DS: Solution to Moderately Sophisticated Social Mining for Elicitation G
DC: Mitigate Moderately Sophisticated Social Mining for Elicitation R
OC: Effectively uses moderately sophisticated techniques to recruit persons for espionage and sabot...
DS: Mitigate Moderately Sophisticated Recruitment Techniques
DC: Mitigate Moderately Sophisticated Recruitment Techniques
OC: Effectively uses highly sophisticated techniques to recruit persons for espionage and sabotage
DS: Mitigate Highly Sophisticated Recruitment Techniques
DC: Mitigate Highly Sophisticated Recruitment Techniques
OC: Effectively uses rudimentary social engineering attacks
DS: Solution to Rudimentary Social Engineering R
DC: Mitigate Rudimentary Social Engineering R
DS: Solution to Rudimentary Social Engineering Full GR
DC: Mitigate Rudimentary Social Engineering GR
DS: Solution to Rudimentary Social Engineering G
DC: Mitigate Rudimentary Social Engineering G
OC: Effectively uses moderately sophisticated social engineering attacks
DS: Solution to Moderately Sophisticated Social Engineering G
DC: Mitigate Moderately Sophisticated Social Engineering G
DS: Solution to Moderately Sophisticated Social Engineering R
DC: Mitigate Moderately Sophisticated Social Engineering R
DS: Solution to Moderately Sophisticated Social Engineering Full GR
DC: Mitigate Moderately Sophisticated Social Engineering GR

ASSET-TYPE: Computing Device
OC: Use compromised humans to attack data on devices and links through highly-sophisticated social engineering/elicitat...
DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated social engineer...
DC: Mitigate Highly Sophisticated Social Mining for Elicitation G
DC: Mitigate Highly Sophisticated Social Engineering GR
DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated social engineer...
DC: Mitigate Highly Sophisticated Social Mining for Elicitation R
DC: Mitigate Highly Sophisticated Social Engineering GR

- DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated social engineer...
  - DC: Mitigate Highly Sophisticated Social Mining for Elicitation GR
  - DC: Mitigate Highly Sophisticated Social Engineering GR

OC: Exploit Known and Unknown Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
- DS: BS: Detect and Respond to exploitation of Known Weak Configurations Settings (CCEs) in Software (OS, firmware, Appl...
  - DC: BG: Protect against Known Weak Configurations Settings (CCEs)

OC: Use compromised humans to attack data on devices and links through basic recruitment through moderately-sophisticat...
- DS: Mitigate use of compromised humans to attack data on devices and links through basic recruitment through moderately...
  - DC: Mitigate Moderately Sophisticated Recruitment Techniques

OC: Use compromised humans to attack data on devices and links through rudimentary social engineering
- DS: Mitigate use of compromised humans to attack data on devices and links through rudimentary social engineering G
  - DC: Mitigate Rudimentary Social Engineering G
- DS: Mitigate use of compromised humans to attack data on devices and links through rudimentary social engineering R
  - DC: Mitigate Rudimentary Social Engineering R
- DS: Mitigate use of compromised humans to attack data on devices and links through rudimentary social engineering GR
  - DC: Mitigate Rudimentary Social Engineering GR

OC: Exploit Known and Unknown Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
- DS: BS: Protect against Known Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
  - DC: BG: Protect against Known Weak Configurations Settings (CCEs)

OC: Use compromised humans to attack data on devices and links through moderately-sophisticated social engineering/elic...
- DS: Mitigate use of compromised humans to attack data on devices and links through moderately-sophisticated social engi...
  - DC: Mitigate Moderately Sophisticated Social Engineering GR
  - DC: Mitigate Moderately Sophisticated Social Mining for Elicitation GR
- DS: Mitigate use of compromised humans to attack data on devices and links through moderately-sophisticated social engi...
  - DC: Mitigate Moderately Sophisticated Social Engineering G
  - DC: Mitigate Moderately Sophisticated Social Mining for Elicitation G
- DS: Mitigate use of compromised humans to attack data on devices and links through moderately-sophisticated social engi...
  - DC: Mitigate Moderately Sophisticated Social Engineering R
  - DC: Mitigate Moderately Sophisticated Social Mining for Elicitation R

OC: Exploit Known and Unknown Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor) o...
- DS: BS: Limit damage from Known Weak Configurations Settings (CCEs) in Software (OS, firmware, Application, Hypervisor)...
  - DC: BG: Protect against Known Weak Configurations Settings (CCEs)

OC: Can compromise data on computing devices, wired links and cryptographic mechanisms that are unprotected
- DS: Can mitigate attacks on data on RF Links that are unprotected using cryptography
  - DC: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
- DS: Can mitigate attacks on data on RF Links that are unprotected using cryptography and physical protections
  - DC: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...
  - DC: Mitigate ability to defeat a weak commercial cryptographic mechanism in a computing device
- DS: Can mitigate attacks on data on RF Links that are unprotected using physical access controls
  - DC: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...

OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with simple protections
- DS: Can mitigate attacks on data on RF Links with simple protections using physical access controls
  - DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 2, or less) with ...

OC: Use compromised humans to attack data on devices and links through highly-sophisticated recruitment for espionage/...
- DS: Mitigate use of compromised humans to attack data on devices and links through highly-sophisticated recruitment fo...
  - DC: Mitigate Highly Sophisticated Recruitment Techniques

OC: Inject Hardware
- DS: Mitigate hardware injection
  - DC: Mitigate Hardware Injection

OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with highly-sophisticated protec...
- DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using cryptography
  - DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material
- DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using cryptography and physical prot...
  - DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection, classifi...
  - DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)
- DS: Can mitigate attacks on data on RF Links with highly-sophisticated protections using physical access controls
  - DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...

OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with extra highly-sophisticated ...
- DS: Can mitigate attacks on data on RF Links with extra highly-sophisticated protections using physical access controls...
  - DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
- DS: Can mitigate attacks on data on RF Links with extra highly-sophisticated protections using cryptography
  - DC: Mitigate ability to defeat Military Grade Crypto by obtaining key material (faster than T5)

OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with rudimentary protections
- DS: Can mitigate attacks on data on RF Links with rudimentary protections using physical access controls
  - DC: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
- DS: Can mitigate attacks on data on RF Links with rudimentary protections using cryptography

DC: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
DS: Can mitigate attacks on data on RF Links with simple protections using cryptography and physical protections
DC: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
DC: Mitigate ability to defeat Strong Commercial Crypto by obtaining key material
OC: Exploit Hardware Vulnerabilities
DS: Mitigate hardware vulnerability
DC: Mitigate Vulnerable Hardware
OC: Can compromise data on computing devices, wired links and cryptographic mechanisms with moderately-sophisticated pr...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography reducing vuln...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls a...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls b...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using physical access controls c...
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography using stronge...
DC: Mitigate ability to defeat Commercial Crypto using military grade encryption
DS: Can mitigate attacks on data on RF Links with rudimentary protections using cryptography and physical protections a...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis
DS: Can mitigate attacks on data on RF Links with moderately-sophisticated protections using cryptography and physical ...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DC: Mitigate ability to defeat Commercial Crypto using cryptanalysis

ASSET-TYPE: Physical Space
OC: Can obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DS: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
DC: Mitigate ability to obtain physical access to classified systems in SCIFs (Protection 5) with minimal stealth
OC: Can obtain physical access to classified systems with light or heavy physical protection (protection 3-4) with mod...
DS: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection (protecti...
OC: Can obtain physical access to classified systems with light physical protection (Protection 3 or less) with minim...
DS: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
DC: Mitigate ability to obtain physical access to classified systems with light physical protection (Protection 3 or ...
OC: Can obtain physical access to cryptographic mechanisms and keys (Protection 2.5, or less) with no stealth
DS: Mitigate ability to obtain physical access to cryptographic mechanisms and keys (Protection 2.5, or less) with no ...
DC: Mitigate ability to obtain physical access to cryptographic mechanisms and keys (Protection 2.5, or less) with no ...
OC: Can obtain physical access to access-controlled unclassified systems (Protection 1-2) with high stealth
DS: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 1-2) with high s...
OC: Can obtain physical access to classified systems with light or heavy physical protection, classified systems in SC...
DS: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection, classifi...
DC: Mitigate ability to obtain physical access to classified systems with light or heavy physical protection, classifi...
OC: Can obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal stealth.
DS: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...
DC: Mitigate ability to obtain physical access to poorly-protected, unclassified systems (Protection 1) with minimal st...
OC: Can obtain physical access to access-controlled unclassified systems (Protection 2, or less) with moderate stealth...
DS: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 2, or less) with ...
DC: Mitigate ability to obtain physical access to access-controlled unclassified systems (Protection 2, or less) with ...

Appendix B - Additional Information on EVRA

This appendix provides additional information on EVRA. The appendix is divided into two sections: (1) a brief summary of the EVRA methodology and (2) detailed results and timekeeping data for EVRA during its application on the comparative study.

Summary of EVRA Methodology

EVRA analyzes (1) a set of mission/business objectives that depend on a cyber system, (2) cyber threats that could impact mission/business objectives by attacking the underlying system, and (3) details of the cyber system upon which the mission/business objectives depend. The process is intended to help answer three key questions:

(1) If a threat action was carried out, what would be the mission impact be?

(2) What adversary level of capability (LOC) is required, as estimated along the DSB scale from I to VI (Gosler & Von Thaer, 2013)?

(3) What mitigation options for are available to deal with the threats, particularly those that have low LOC and high mission impact?

The EVRA processes (Figure 29) maps well to the NIST risk assessment framework (National Institute of Standards and Technology, 2012). One difference is that whereas EVRA uses LOC, NIST uses likelihood of attack. We believe, but lack the empirical data to strongly support, that LOC correlates to likelihood of attack. The reasoning is that (1) attacker motivation is assumed and (2) by possessing sufficient LOC, attack likelihood goes up.

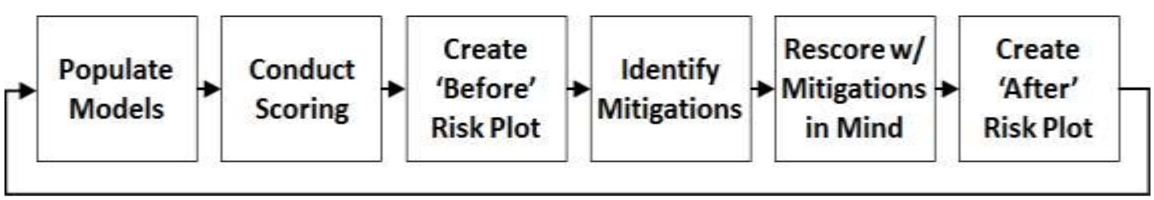

Figure 29: Summary of the EVRA Methodology

An overview of Figure 29 is as follows. First, SAs obtain data to populate the adversary, mission, and system models. Next, SAs score a set of potential attacks and estimate risk using the models. Scoring is along a 5-point Likert-style ordinal scale. SAs score attack LOC and mission impact if the attack is successful. Then the scoring data is entered into a tool which produces the initial EVRA risk plot. After deciding on possible mitigations and rescoring LOC scores appropriately given the hypothetical presence of the mitigations, the SAs rerun the EVRA tool to obtain an "after" risk plot. The results are then shared with other stakeholders for decisions on the way forward. As threat, mission, and system change over time, the entire process iterates.

Omega Data Capture and Timekeeping Data

This section presents detailed results and timekeeping data for the EVRA team. Table 23 shows the scoring table the EVRA team used to record starting node LOC scores and the corresponding rationale the team recoded for each score.

Table 23: EVRA Starting LOC Scoring and Rationale

Vector	Node	Starting LOC Score	Starting LOC Rationale
Malicious Insider	Storage Server	2	Easy for someone with access to a fileserver to upload a malicious payload. 2 not 1 b/c requires stealth to bypass detection. 2 not 3 because person is an insider and may understand what protections are in place and how to bypass them.
	Apps/tools	2	Need to modify application to ultimately cause bad effects downstream - maybe SE admin to give greater permissions for modifying config files if needed/don't already have
	Linux OS		Assuming user is familiar with Linux, same as above
	HP WS HW	2	HW implant - user just plugs something in that they did not have to develop, but needs to be somwhat discreet
	Ground Segment Network	2	User has access to switch management console, just has to be careful to not get caught by other admins
	Authentication Service	2	User has access to DC, just has to be careful to not get caught by other admins
	Data Switch	2	be making change to config to pass data to unintended workstation
	Crypto	1	Assuming start/target are same, and any attack just translates to availability (zeroize or change key)
	Premise Router	2	User has access to management console, just has to be careful to not get caught by other admins - attack might be making change to config to pass data to unintended workstation - external connectivity to this router is out of scope
Unwitting Insider	Storage Server	3	Example: user takes file home each night and uploads to file server each morning, adversary knows this after some
	Apps/tools	3	Example: Configuration of apps is dictated by unclass documentation (also goes for admin tools like STIG lists) - adversary makes changes to these unclass templates or documentation, causing SA to make incorrect updates allowing future access - needs to be stealthy
	Linux OS	3	Bad guy poisons repo typically used for OS updates, maybe that gets pulled into local repo
	HP WS HW	4	Trick user into installing HW implant, maybe hidden in expensive peripheral device that they "won" - score mostly based on level of difficulty for developing device
	Ground Segment Network	3	Patch/update for network device - requires background research for device details, as well as spoofing email/vendor website post, and potentially forge digitial signature to get a bad update installed
	Authentication Service	3	Example: adversary already has access to a low level account, but cant escalate priv. Send spoofed email to SA to grant additional priv.
	Data Switch	3	Patch/update for network device - requires background research for device details, as well as spoofing email/vendor
	Crypto	3	Either send spoofed email telling user to zerize crypto, or send bad crypto to be re-loaded - could also send a bunch of garbage to make them think crypto load is bad
	Premise Router	3	Patch/update for network device - requires background research for device details, as well as spoofing email/vendor website post, and potentially forge digitial signature - external connectivity to this router is out of scope
Network	Premise Router	5	Get access to management console and make configuration changes to enable other attacks - Tier 5 is the first level that has ability to inject traffic into classified network
	Data Switch	5	Get access to management console and make configuration changes to enable other attacks - Tier 5 is the first level that has ability to inject traffic into classified network
Supply Chain	Storage Server	4	No need to infiltrate manufacturing - background research to determine what is being used in facility, and background research to get UPS driver to drop off compromised device
	Apps/tools	5	Must be able to infiltrate SW development site, and develop high-stealth implant
	Linux OS	5	Must be able to infiltrate SW development site, as develop high-stealth implant
	HP WS HW	4	No need to infiltrate manufacturing - background research to determine what is being used in facility, and background research to get UPS driver to drop off compromised device
	Ground Segment Network	4	No need to infiltrate manufacturing - background research to determine what is being used in facility, and background research to get UPS driver to drop off compromised device
	Authentication Service	4	No need to infiltrate manufacturing - background research to determine what is being used in facility, and background research to get UPS driver to drop off compromised device
	Data Switch	4	No need to infiltrate manufacturing - background research to determine what is being used in facility, and background research to get UPS driver to drop off compromised device
	Crypto	5	Need to either intercept keying material or modifiy encrypter before delivery
	Premise Router	4	No need to infiltrate manufacturing - background research to determine what is being used in facility, and background research to get UPS driver to drop off compromised device

Table 24 shows the shows the scoring table the EVRA team used to record target node LOC scores and the corresponding rationale for each score.

Table 24: Target LOC Scores and Rationale

Node	Target LOC (C)	Target LOC Rationale (C)
Storage Server	1	Assuming domain controller is compromised, attacker could change permissions to see all network shares and their content.
Apps/tools	1	Assuming applications do not encrypt or obfuscate their data and internal communications, and it's written to file system
Linux OS	1	Assuming domain controller is compromised, attacker has access all data running on workstation
HP WS HW	N/A?	NA?
Ground Segment Network	2	Assuming domain controller is compromised and switch access is centrally managed by DC, remotely modify switch configs to route traffic to unintended recipients - requires networking expertise
Authentication Service	1	Assuming admin access on networked workstation, remote to DC and modify files
Data Switch	2	Assuming domain controller is compromised and switch access is centrally managed by DC, remotely modify switch configs to route traffic to unintended recipients - assumes attack is on red switch - requires networking expertise
Crypto	N/A?	Attack would require physical access? Not able to compromise via network?
Premise Router	2	Assuming ground segment switch is compromised, route traffic from SIPR router to unintended recipients- some network expertise required

Node	Target LOC (I)	Target LOC Rationale (I)
Storage Server	1	Assuming domain controller is compromised, attacker could change permissions to modify contents of storage server
Apps/tools	3	Assuming real-time incoming data stream and need to intercept and modify. Also assuming relatively easy access to documentation describing how the apps receive and process their data.
Linux OS	1	Assuming domain controller is compromised, attacker has access all data running on workstation
HP WS HW	N/A?	NA?
Ground Segment Network	2	Assuming system on management network is compromised and switch access is centrally managed by DC, modify config to ignore/drop critical protocols - requires networking expertise
Authentication Service	1	Assuming admin access on networked workstation, repeatedly reboot DC
Data Switch	2	Assuming system on management network is compromised and switch access is centrally managed by DC, modify config to ignore/drop critical protocols - assumes attack is on red router - requires
Crypto	N/A?	Attack would require physical access? Not able to compromise via network?
Premise Router	2	Assuming ground segment switch is compromised, drop all traffic coming from SIPR router - some network expertise required

Node	Target LOC (A)	Target LOC Rationale (A)
Storage Server	1	Assuming domain controller is compromised, fill up audit logs until server crashes.
Apps/tools	1	Assuming domain controller is compromised, change permissions so users can not launch application, or uninstall it
Linux OS	1	Assuming domain controller is compromised, attacker has access all data running on workstation
HP WS HW	1	Assuming domain controller is compromised, attacker can deny access to all users
Ground Segment Network	2	Assuming system on management network is compromised and switch access is centrally managed by DC, modify config file to shut ports and change root password, or upload bad boot config - requires networking expertise
Authentication Service	1	Assuming admin access on networked workstation, repeatedly reboot DC
Data Switch	2	Assuming system on management network is compromised and switch access is centrally managed by DC, modify config file to shut ports and change root password, or upload bad boot config - assumes attack is on red router - requires networking expertise
Crypto	N/A?	Attack would require physical access? Not able to compromise via network?
Premise Router	1	Assuming ground segment switch is compromised, flood SIPR router with traffic (ping, SYN/ACKs)

Notes kept by the EVRA team during their analysis of Omega are given in Figure 30.

Note: Additional logic captured as rationale in scoring spreadsheet.

8/21/17 – LOC scoring

Assumption:

- Path LOCs not considered (1 – never hardest part of an attack)
- Not considering different roles – assuming compromised user is an admin (Worst case)
- Ground Segment Network is a layer 2 switch
- Equal difficulty to compromise all data types
- All data within scope is unencrypted
- Vector not relevant for Target LOC scoring
- Starting LOC scores do not consider development of payload
- Apps are all COTS from contractor (SCADA/ICS-like)
- Linux OS updates are not hashed/checked before installation
- Looks like mitigations prevent most attacks by tier 1
- Going with worst case – malicious insider is privileged user
- Assuming separation of duties for privileged users
- Malicious insider: Only one person has turned and is malicious
- Authentication service is a DC
- Data switch attack is on red switch, left of crypto
- Threat facing the unwitting insider is external to the organization
- Due to existing mitigations, unwitting insider isn't an admin that misconfigured something – also external to the organization, otherwise, it's a malicious insider
- Workstations do not leave secure facility
- Satellite is sending classified data to GEP segment/Data switch
- Using an off the shelf Linux distro
- OS and Apps are updated regularly
- Policy best practices not implemented (separation of duties, job rotation)

Figure 30: Notes from EVRA team

Table 25 contains the total hours by day for the EVRA team over the seven-day period in which they conducted the EVRA analysis of Omega. As shown, the team spent a total of 24.95 hours on the task, with 14.30 hours spent on choosing and scoring attacks for the "before" risk plot (BP) and 10.65 hours on selecting mitigations based on the "before" plot and scoring LOC assuming the mitigations are in place.

Day	Total Hours	# People	Category	Description of Activity
Day 1	0.60	1	BP-Choose	Preparing for LOC scoring session/building scoring worksheet
Day 4	0.30	1	BP-Choose	Preparing for LOC scoring session/building scoring worksheet
Day 5	2.00	1	BP-Choose	Preparing for LOC scoring session/building scoring worksheet
Day 5	4.20	3	BP-Score	Start LOC scoring
Day 6	7.20	3	BP-Score	More Start LOC scoring, plus Target LOC scoring
Subtotal	**14.30**			
Day 6	1.50	3	AM-Cut Line	Discussion on where the cut line is - mitigations need to eliminate 1's and 2's
Day 6	3.75	3	AM-Select Mitigations	Develop mitigations to reduce the Starting and Target LOCs to 3's and higher
Day 7	5.40	3	AP-Score	Start and Target LOC mitigated scoring
Subtotal	**10.65**			
Total	**24.95**			

Table 25: Timekeeping for EVRA Team

Values in the Category column of Table 25 have meanings defined in Table 26, as recorded by the EVRA team.

BP-Choose	Choose attacks to score (T_BP-A)
BP-Score	Score attacks for level of effort (LOE) + rationale (T)_BP-B)
BP-Map	Map mission impact score to each attack (T)_BP-C)
AM-Cut Line	Time to determine a "cut line" to ID highest risk attacks that need mitigation (T_AM-A)
AM-Select Miti	Time to select mitigations for attacks to mitigate (T_AM-B)
AP-Score	Re-score attack LOEs based on addition of mitigations (T_AP-A)

Table 26: Timekeeping Categories for EVRA Analysis