2011

The Correlates of Cyber Warfare: A database for the modern era

Charles Debeck
Iowa State University

The Correlates of Cyber Warfare: a database for the modern era

by

Charles DeBeck

A thesis submitted to the graduate faculty

in partial fulfillment of the requirements for the degrees of

MASTER OF ARTS

MASTER OF SCIENCE

Co-majors: Political Science; Information Assurance

Program of Study Committee:
Steffen Schmidt, Co-major Professor
Thomas Daniels, Co-major Professor
Barbara Licklider

Iowa State University

Ames, Iowa

2011

Dedication

I would like to dedicate this thesis to my family who has continued to support me throughout my educational career and my many friends and colleagues who have continued to support me academically, emotionally, and intellectually while I pursued this thesis. I would also like to dedicate this thesis to the many important figures throughout my life who have allowed me to get where I am, including but by no means limited to Dr. Hoekstra, Professor Karen Cherewatuk, and Ms. Heidi Gusa.

Table of Contents

Acknowledgements

I would like to thank my major professors, Dr. Thomas Daniels and Dr. Steffen Schmidt, for their support throughout the thesis writing process. Professor Barbara Licklider and Senior Lecturer Jan Wiersema have also been extremely helpful throughout my graduate career and deserve acknowledgement. I would like to thank the National Science Foundation for sponsoring the Scholarship for Service program which I have had the honor in participating in through my graduate career. I would like to thank Dr. Doug Jacobson for his technical expertise and general assistance. I would like to thank Virginia (Ginny) Anderson for her mastery of all things in the information assurance department. I would like to thank the staff of the Iowa State University political science department for their assistance throughout my educational career. I would like to thank the many wonderful professors I had during my undergraduate career at St. Olaf College who are too numerous to name. I would like to thank Professor Anthony Lott for his direct assistance with my thesis formulation and correspondence throughout the creation process.

Abstract

The author proposes a new theoretical database for tracking cyber attacks on the internet called the “Correlates of Cyber Warfare”. Inspired by Professor J. David Singer’s Correlates of War project, the author provides an overview of the requirements for such a database to be feasibly implemented and the benefits such a database would provide for future research. The author begins by covering the various databases which exist to track and monitor online disturbances, previous research involving collaborative intrusion detection systems, and the lack of data available regarding the sheer volume of cyber attacks traversing the internet. The paper continues to outline what data could be collected by a strategic deployment of IDS’s, methods for optimal IDS deployment (with emphasis on transnational router collection), and potential privacy pitfalls that are summarily addressed. The paper ends by describing the many research and online safety benefits that would be seen by the implementation of such a system and methods to work around incomplete implementation.

Section Zero: Introduction

War is a complicated subject. The Correlates of War Project (COW), created by Professor J. David Singer in 1963, was made to be used as a database from which correlates of the likelihood of war could be drawn. The COW data now is used by political scientists everywhere and often is drawn upon to inform policies and policymakers related to international relations.

What about cyber war? While the correlates of war are considered very interesting, there is startlingly little data about the correlates of cyber war. What makes a state more or less likely to engage in cyber attacks? What countries are used as jump points for cyber attacks? Which countries are the targets of the most cyber attacks per capita? These and other interesting questions cannot be answered with the current data available to scholars.

This paper intends to provide an outline, a proof of concept, of a database which could help answer these questions, the Correlates of Cyber Warfare (COCW). The COCW is a theoretical database being proposed by the author as a new database to measure cyber attacks on the internet. The COCW would provide the exact number of cyber attacks passing through certain points on the internet and the nature of these attacks. The paper will begin in Section One by defining the term "cyber attack" and what definition of "cyber attack" the COCW database will use as well as explaining what an Intrusion Detection System (IDS) is. Section Two will continue by outlining what databases currently exist in regards to cyber attacks and what unique and different benefits the COCW will provide. Section Three will go into a brief literature review discussing similar research done in collaborative Intrusion Detection System (IDS) environments and how this research relates to the underlying principles of the COCW database. In Section Four the author will discuss what data will

need to be collected for the COCW database to be effective, the key variables the COCW will display upon implementation, and the importance of these variables. Section Five will be a discussion of different methods by which to collect data, the benefits of each method, and ultimately a recommendation for which method to use. Section Six will discuss the major privacy issues surrounding the COCW system and how the author intends to address such issues. Section Seven will delve into the technical details and recommendations on data collection strategies using recent research into deep packet inspection at high traffic volumes, data sanitization techniques, and symmetric key encryption. Section Eight will describe some of the broader implications of implementing the COCW system on a national or international scale, such as increased international cyber security, increased domestic cyber security, and opening new avenues of research for scientists globally. Section Nine will address the realistic possibility of incomplete implementation solutions for the COCW. The author will then conclude with parting thoughts and recommendations for future avenues of research.

Section One: Definitions

Section 1-1: What is a Cyber Attack?

In order to understand the Correlates of Cyber Warfare, one must begin by defining what a "cyber attack" (the unit of measure for the database) is. "Cyber attacks" are usually defined by their medium: they are attacks which rely upon the use of information technology as the medium of attack. This conceptualization does not preclude the final target need be a computer; in fact, many attacks have very tangible, non-digital consequences. To understand what cyber attacks are, this paper will go through a number of different types of attacks showing the diverse array of attacks available to a malicious or simply ignorant actor.[1]

The quintessential network attack is a denial of service attack (DOS). A DOS attack tries to access one site, server, or point of the internet over and over again in quick succession. The target, being flooded with requests to connect or access, quickly fills up any available space it has for incoming access. Suddenly the targeted point can no longer accept legitimate requests from users, rendering the target unavailable and effectively crashed. The more complex (and generally more common) counterpart is the distributed denial of service attack (DDOS) which uses multiple computers to try and access the targeted point, thereby reducing the load on each individual attacking computer and increasing the volume of attack on the target.

A classic attack that has infected millions of computers over the years is the Trojan Horse. Named after the Iliad's famed icon, this attack convinces the user of a computer to download an infected file. The file, once inside the walls of the computer, can cause all sorts

[1] The following examples are drawn from the website of ESET North America, an internet security company headquartered in San Diego: <http://securingourecity.org/>

of chaos, from deleting important files to allowing the attacker remote access into your computer. A similar attack, a virus, requires the user to in some manner allow the virus access to the computer (usually unknowingly) before unloading its destructive payload.[2]

The final attack is also the most well-known: the computer worm. Worms traditionally have one main purpose, to propagate. A computer worm requires no action on the part of the user, it simply works within the computer to create and send out copies of itself to as many connected systems as it can find. Such an attack can cause the same traffic issues discussed with the DOS attacks, making worms a dangerous and very real threat.

These attacks and many more are found every day by security experts. Due to the diversity of these attacks, there exist large databases of signature files which describe individual attacks and their characteristics (usually containing bits of code found in the attack). Recognizing this diversity, for the purposes of the COCW database a cyber attack will be defined as "any attack which can be detected by a given Intrusion Detection System (defined in Section 1-2) using up-to-date cyber attack signatures". Attacks which fall into this category include but are not limited to worms, spam, Trojan horses, viruses, and denial of service attacks. The tautological nature of this definition should be noticed; if the attack can be detected, then it is undeniably an attack, but if the attack is not detected for the purposes of this database there was no attack. As signature files get better the number of attacks will naturally increase due to increased detection rates. The author uses this self-fulfilling definition to assure 100% cyber attack detection for calculation purposes but recognizes the limits of such a tautology. However, the author feels the benefits gained from

[2] *Ibid.*

limiting the definition of "cyber attack" far outweigh the costs for this database and as such moves forward with this understanding. This tautology is further addressed in Section 9-1.

Section 1-2: What is an "Intrusion Detection System"?

An Intrusion Detection System (IDS) is a software or hardware system which detects attacks from outside parties. Many individual networks use IDS's to mark packets which have potentially harmful qualities. These qualities are often discerned using "deep packet inspection". This form of inspection looks into the contents of each message passed through the IDS for a number of "signatures", known malicious code, to determine if the message is malicious in nature. Such inspection can be computationally intensive. A deep packet inspecting IDS typically looks for malware, spam, viruses, and other cyber attacks.[3] When an IDS detects a threatening packet the system logs the packet information along with which signature was violated.[4] These logs include a host of relevant information, such as the source and destination IP address, and the source and destination hardware address, and the date and time of detection.[5] An IDS can be deployed "on top" of a network or "in line" with a network. For the purposes of the COCW, the IDS will be in line, meaning that any data being sent over a wire will pass into the IDS, be inspected, and then be passed out of the IDS and on to its destination. For the purposes of the COCW, an IDS will be defined as "an intrusion detection system using deep packet inspection and a reasonably updated signature database in line with, not on top of, network flow."

---

[3] The SANS Institute Reading Room (2001). Understanding Intrusion Detection Systems. <http://www.sans.org/reading_room/whitepapers/detection/understanding-intrusion-detection-systems_337>
[4] *Ibid.*
[5] *Ibid.*

The IDS will be a fundamental component in the design and implementation of the COCW database. Section Five will outline a number of collection methods which utilize the convenience of IDS's to collect the necessary data for the COCW database to function.

Section Two: Current Databases

This section will discuss the currently existing databases categorizing cyber attacks found in the wild. Specifically, the section will look at three databases: the Common Vulnerabilities and Exposure listings which provide comprehensive descriptions of known vulnerabilities and exposures, the United States Computer Emergency Readiness Team's (US-CERT) Vulnerability Notes database which provides the most up to date knowledge of vulnerabilities, and the Internet Storm Center (ISC) which provides predictions on potential attack vectors. Each database provides a unique contribution to the security community, yet none of the databases provide the information the COCW would provide. The section will outline how each database functions, the good ideas and bad ideas of each database, and ultimately how the database relates to the COCW.

Section 2-1: Mitre Corporation's Common Vulnerabilities and Exposures

The Common Vulnerability and Exposures (CVE) is a list of vulnerability definitions to be used by the security community. The purpose of the CVE is to provide a common language for security professionals to use when discussing a given vulnerability.[6] Some of the objectives of the CVE include only assigning a single name to any given vulnerability or exposure, providing a single description for each vulnerability or exposure, and to provide the capacity for interoperability between databases and security resources.[7]

The CVE imports new vulnerabilities and exposures via independent submission coupled with rigorous review of data. If a vulnerability or exposure is found, then it is

[6] Common Vulnerabilities and Exposures (2010). About CVE. < http://cve.mitre.org/about/index.html>.
[7] *Ibid.*

assigned a "candidate" status by a "CVE Candidate Naming Authority (CNA)".[8] The candidate is then put before the CVE Editorial Board which votes as to whether or not the proposed candidate becomes a CVE entry. If it does not pass, the reasons for its refusal are posted; however, if accepted the candidate receives the designation "entry" in the CVE List. Through this rigorous methodology CVE listed vulnerabilities and exposures gain entrance into the CVE listing.

The rigor of the CVE, though providing some ideal qualities in a database, lacks a practicality for the level of traffic entering the COCW. The COCW database should in all likelihood be receiving hundreds if not thousands of entries daily; the rigorous method of the CVE would slow the admission process to a crawl. As such, the use of a voting committee would be impractical, but the lesson of the CVE can still be incorporated into the COCW. Entries into the COCW database need to be validated in some manner, such as the use of encrypted submissions as discussed in Section 7-3, and the COCW could even draw upon the definitions provided by the CVE to define the attacks that have been detected. Though the CVE provides a very useful common language for security professionals, it cannot provide sufficient information to engage in empirical research like the COCW would enable.

Section 2-2: The US-CERT Vulnerability Notes Database

The United States Computer Emergency Readiness Team (US-CERT) database publishes information about vulnerabilities submitted by individuals to the US-CERT. This database allows for the publishing of potentially less-severe vulnerabilities.[9] The US-CERT database has the ability to search through vulnerabilities using a keyword search to find

[8] *Ibid.*
[9] United States Computer Emergency Readiness Team (2011). US-CERT Vulnerability Notes. <http://www.kb.cert.org/vuls/>

specific vulnerability information of desired. Individual administrators can report discovered vulnerabilities, find recently published vulnerabilities, and search for vulnerabilities that might be relevant to their systems.[10] Independent action plays a large role in the US-CERT database system.

Vulnerabilities are reported to US-CERT using two methods: online form or e-mail. The online form does not require identity authentication by the submitter.[11] Likewise, the e-mail option does not require authentication, an individual simply e-mails the information they have available to the database. There are no specific requirements for submission of vulnerability information; oftentimes discoverers simply submit as much information as they can. The anonymity is meant to allow for greater submission rates without fear of release of information; the freedom in submission style is meant to create the potential for incidental discovery.[12]

The US-CERT database lacks a particular statistical rigor preventing researchers from engaging in the analysis which the COCW aims to provide. First, the US-CERT Vulnerability Notes database relies on individual submission of vulnerability information which inevitably raises selection bias concerns. Furthermore, there is a lack of standards in the submission of information. Two individuals observing the same vulnerability may report the vulnerability completely differently with the first individual submitting multitudes of details related to the vulnerability and the latter individual sending in just the minimal information. Such gaps in reporting can lead to vulnerabilities being reported multiple times or different vulnerabilities not being sufficiently differentiated. Finally, the US-CERT does

[10] *Ibid*.
[11] US-CERT (2011). Vulnerability Reporting Form. <https://forms.cert.org/VulReport/>.
[12] *Ibid*.

not authenticate the actual individuals submitting information. If a malicious individual wished to poison the database with false or misleading information the challenge to doing so would be relatively low with almost no chance of the false lead being attributed to the attacker. However, the US-CERT does have an advantage in its ability to report vulnerabilities with relative alacrity; unlike the CVE which requires time to process the US-CERT database can output information very quickly. This advantage will be important in the COCW and is an ideal to be aimed for without succumbing to the potential pitfalls outlined above.

Section 2-3: The SANS Institute's Internet Storm Center (ISC)

The Internet Storm Center is a database hosted by the SysAdmin, Audit, Network, Security (SANS) Institute. Established in 2001, the ISC aims to provide "free analysis and warning service to thousands of Internet users and organizations".[13] The database uses collections of individually submitted logs of security events to detect correlations and potentially new incoming attacks before the spread of the attack grows too far. This correlation is fueled by the aggregation of IDS logs from around the world volunteered by different organizations and users.[14]

The submission of these logs follows a common, automated format which can be easily parsed into the relevant data at the ISC servers. The ISC provides a patch for many common IDS's and firewalls to automatically send information in this defined format.[15] One goal of the ISC collection method is to mirror the techniques used by weather forecasters; by having large amounts of global data the ISC can look for trends and potential upcoming

[13] The SANS Institute (2011). About the Internet Storm Center. <http://isc.sans.edu/about.html>.
[14] *Ibid.*
[15] *Ibid.*

attack methods.[16] As such, anyone can participate in the ISC by implementing the publicly provided software.

The ISC provides an interesting glimpse of what the COCW could be if it was forward looking. The consistent data submission format in the ISC adds a convenience to data processing not previously seen in the other databases presented. Furthermore, the ISC has shown that large enough amounts of data can provide actionable intelligence and allow for the prevention of attacks even before they occur. However, the ISC lacks a statistical rigor in its open-ended implementation and risks a selection bias. Furthermore, the ISC aims to provide predictive data of what will happen within the next 24 to 48 hours. The COCW is a past-oriented database, looking to analyze what has already happened. The software and ideas of data submission can definitely be gained by studying the ISC, but there is no doubt that the contributions of the COCW database are unique from those of the ISC.

Section 2-4: Conclusions

While these databases provide important syntax and material for information security professionals, they lack statistical data on the numbers of cyber attacks occurring in the wild. The COCW would provide something wholly unique: statistically viable information on the number of cyber attacks occurring, where they are coming from, where they are going, and a host of other useful information described in Section Four. While the above databases are undoubtedly useful tools, the need for a new database with actual numbers of cyber attacks would prove invaluable for researchers and security professionals internationally. For these reasons the author proposes the creation of the COCW database.

---

[16] *Ibid.*

Section Three: Literature Review

Previous literature has discussed the concept of a "cooperative IDS system" which would be used for more accurate detection of attempted attacks. By reviewing this literature the COCW will glean the benefits of IDS systems working together, some potential drawbacks that will need to be addressed, and some implementation schemes that have been tried before.

The concept of a cooperative intrusion detection system rests on the use of multiple intrusion detection systems to better detect and classify attacks on a given network. For an explanation of what intrusion detection systems are reference Section 1-2. These cooperative systems were drawn out of research discovering that complex attacks on a single network would not be detected by a single IDS but in fact multiple IDS's working in tandem could detect the attempted breach.[17] These sorts of attacks became increasingly relevant as researchers discovered that attackers were attacking from multiple distinct systems at once rather than simply one computer.[18]

Early cooperative systems relied on a central database or server from which commands could be given, information could be compiled, and decisions could be made.[19] These systems outlined a number of issues which would be present in a cooperative, data-sharing solution. First, there must be an aspect of "local trust".[20] This trust requires that

---

[17] Porras P. and P. Neumann (1997). EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances. *Proceedings of the 1997 National Computer Security Systems Conference*, pages 719-729. Drawn from Frincke, Deborah et al. (1998). A Framework for Cooperative Intrusion Detection. <http://csrc.nist.gov/nissc/1998/proceedings/paperF6.pdf>.

[18] Snapp S., J. Brentano, et al. (1991). DIDS (Distributed Intrusion Detection System) – Motivation, Architecture and an Early Prototype. Drawn from Frincke, Deborah et al. (1998). Page 1.

[19] Frincke, Deborah (2000). Balancing Cooperation and Risk in Intrusion Detection. Pages 1-2. <http://www.cs.ubc.ca/~malam/papers/IntrusionDetection/Frincke%20-%20Balancing%20Cooperation%20and%20Risk%20in%20Intrusion%20Detection%202000%20acm.pdf>.

[20] Frincke, Deborah et al. (1998). Page 2.

each IDS must be in some way trustworthy, whether this is through validation of identity, ensuring the IDS has met required standards, or some other trust-producing protocol. Furthermore, these systems separated the enforcement of policy from the identification of policy violations.[21] This fact mirrors the design of the COCW as enforcement of policy violations is left to someone else entirely; the COCW only looks at the identification of policy violations and reports them. A key component of early designs is that "data collectors should [. . .] provide both data reduction and sanitization"[22]. This focus on protecting data between collector and aggregator will be addressed in Section Six. An interesting concept presented in the literature was that local data collection sites would maintain control over their policy decisions, allowing for easier integration with already existing domains and preventing the need to impose policy decisions on unwilling hosts.[23] By allowing for local autonomy, the data collectors may have collected some irrelevant data, but such data would simply be thrown out by the central management site.[24]

These ideas helped define the concept of cooperative intrusion detection systems working together to provide greater security for networks. With these ideas a number of systems were designed, such as Hummer, EMERALD, and APHID.[25] These systems, though beneficial, were never implemented on a scale or with a purpose like the COCW database, but the COCW can certainly gain from the experience of these systems by following the principles discussed above.

---

[21] *Ibid*. Page 3.
[22] *Ibid.*
[23] *Ibid.*
[24] *Ibid*.
[25] Zamboni, D. and E. Spafford (1999) discusses APHID. Porras P. and P. Neumann (1997) discusses HUMMER. Both references drawn from Frincke, Deborah (2000).

Section Four: The Correlates Of Cyber Warfare Data and Variables

This section will outline the different variables the COCW will contain for analysis. The first section will define exactly what data will be collected by the COCW monitoring system. The second section will present the different variables, discussing the nature of the measurement and its relevance and importance to future study. A number of the variables will be drawn from information inferred from the data; such variables will be noted in the descriptions and the method of inference will be covered in Section 7-4. The variables of the COCW include country of origin, country targeted, country used as vector, target sector, date of release/effect, level of sophistication, attack type, number of cyber attacks, successfulness, and cyber aggressiveness. While not all of these variables need to be implemented in the final database, the author presents these variables as good options for the implementer.

Section 4-1: Data Collected

The COCW monitoring system will collect a very specific amount of data from each detected attack. The technical nature of this collection will be explained further in Sections Five and 7-2. The first piece of data will be the type and time of the attack. This information is gathered automatically by any IDS; the IDS will simply state that a certain flag was set off at a certain time. Second, the destination IP address of the attack will be recorded. This piece of data explains where the attack was going and can provide the inferred data described below. Third, the source hardware address of the attack will be taken in by the COCW, allowing for some source determination. Finally, the source IP address will be taken into account, though the author recognizes the limitations of trusting a source IP address (outlined further in Section Five). However, for many types of attacks spoofing an IP address

eliminates the effectiveness of the attack, making recording the IP address a reasonable step to take in studying the nature of these attacks. In total, these four components comprise the actual data that will be taken in: the type and time of the attack, the destination IP of the attack, the source hardware address of the attack, and the source IP of the attack.

Section 4-2: Country of Origin

Country of origin will be a coded measure of where the cyber attack came from, likely a number for simplicity. The number of countries in the world is irrelevant so long as consistency is maintained. One must recognize the limitations of this measure due to the classic attribution problem of cyber attacks (which will be discussed more in depth in Section Five). Using the source IP address collected by the data above, the database will infer the country of origin by referencing an outside database (as will be outlined in detail in Section 7-4). By coding the country of origin, the database will be able to more effectively see where cyber attacks are coming out of even if such a measure does not completely measure the origin of the attack.

Section 4-3: Country Targeted

This measure will be similar to country of origin, a coded definition of where the attack seemed to be targeting. In the case of a non-targeted or non-specific attack the attack would be coded with a value of zero, representing no known target. An example of such a situation would be a worm released just to spread instead of infect a particular system. However, any non-broadcast packet will contain a destination IP that is collected; by referencing an outside database the COCW can infer the destination/target country. Knowing the country targeted will be useful as a measure of which states are most threatened

by cyber attacks and may lead to interesting research questions as to what causes a country to be targeted by cyber attackers.

Section 4-4: Country as Vector

The country as vector measurement will keep track of an attack's route. For example, if an attack left France, went through three routers in Sweden, and ultimately attacked a computer in the Netherlands, the country as vector measure would add three units to the count for Sweden for this attack (endpoints are not counted in this variable). This variable will be inferred from a combination of two data points: the destination IP address and the country in which the attack was detected. If the attack is detected in a country but the country targeted (as measured above) is not the same as the detecting country, then clearly the country which detected the attack is a "country as vector" as the attack is just passing through. The country as vector measurement would be useful as it would allow for researchers to see which countries are being traversed most often for cyber attacks and may allow for interesting analysis of which countries are most susceptible to penetration or use of their country by cyber attackers.

Section 4-5: Target Sector

Target sector will be a coding of 0, 1, or 2 relating to one of three possible target sectors: industry, private individual, or military/government. Industry will be defined as any privately held company or subsidiary physically located within the country whose primary purpose is to engage in business of some sort regardless of for-profit or non-profit status. A private individual will be defined as any system or network not part of a private industry or the government without a primary purpose of engaging in business. The government and

military will be defined as any branch of the public government system not including publicly funded privately run institutions. To measure this variable the data on destination IP would need to be correlated to a targeted sector. To infer this information would require the database to reference an outside source at some point. Such an external database does not currently exist from the author's research and would need to be created; however, such a database would require substantially less work to make than that required to create the COCW database. This measurement would allow for researchers to study who is being attacked and potentially find correlating factors to such targeting. If in the actual implementation more sectors are considered useful then the number of codes would be increased accordingly.

Section 4-6: Date and Time of Detection

The date and time of detection would note when the attack was discovered, giving information as to the length of time between release and detection. The detection time will be a useful measure as it will allow researchers to know how long attacks are taking to find. As the database collects data over years it will be able to give more detailed analysis with these dates and allow for analysis of how often new attacks are released and discovered. This information will be gathered by using the date and time data collected by the IDS.

Section 4-7: Level of Sophistication and Attack Type

The level of sophistication will provide a glimpse of the technical expertise of the attack being used. This measurement will use a wide variety of weighted qualities to produce an artificial level of sophistication number. This number, between zero and five with non-integer amounts allowed, will allow for comparisons in technical sophistication between

attacks. Some aspects that would add to the level of sophistication would be encryption being used in varying levels, the difficulty in preventing the attack, the amount of time from finding the exploit to producing a satisfactory patch or update, or the difficulty in ascertaining the actual author or source. This customized variable is left to the implementer to define fully, but the author notes the importance of such a variable as it would allow for analysis of correlations to increasingly sophisticated cyber attacks. A customized database correlating different attack signatures to different levels of sophistication would need to be created; from this database the variable would be inferred using the type of attack data collected by the IDS.

Attack type will be a nominal variable describing the type of attack being perpetrated. These types will fall into broad categories commonly delineated by intrusion detection systems (IDS) and simply follow a pattern similar to target sector. For example, a virus may be coded as a 0, a worm as a 1, a Trojan horse as a 2, and so forth, allowing for research into which methods of attack are popular and perhaps even depicting which methods are most commonly used by well-equipped or poorly-equipped criminals. This variable would draw on the "attack type" data as well. To assist with determining the level of sophistication the COCW may draw on the CVE database to assist with syntactically measuring the level of sophistication.

Section 4-8: Number of Cyber Attacks

This measurement will simply be a numeric value of the number of cyber attacks carried out throughout an average day, month, quarter, or year for a given time period. Each attack will be categorized as either outgoing or incoming based on the direction of the attack.

This information requires only the data that an attack occurred being sent by the monitoring system. Such a variable will allow for analysis for peaks and valleys in the level of cyber attack activity.

Section 4-9: "Cyber Aggressiveness"

"Cyber aggressiveness" will be a weighted measure looking at two factors: cyber attacks leaving a country and cyber attacks entering a country. The cyber aggressiveness will be calculated by looking at the proportion of cyber attacks entering a country compared to the number of attacks leaving said country. By comparing this ratio, a country with a substantially higher number of attacks leaving the country than entering will be considered "aggressive". Conversely, a higher number of attacks entering a country than leaving will be labeled "defensive". These designations do not in fact speak to the nature of their populace; rather, they speak to whether attacks are directed domestically or internationally. This variable requires no additional data as it represents a composite measurement of other variables already outlined. This variable is also the only variable measured by country rather than attack.

Section Five: Data Collection

This section will cover the primary challenge to creating this database: data collection. Five methods will be presented. Each method will be rigorously analyzed with notes on the strengths and weaknesses of each plan. Then, a recommended method of data collection will be given by the author with rationale. There are a number of methods by which the database could go about collecting data. Through this section the paper will outline a number of methods which could be used in conjunction with each other or separately to provide the data required to make the database effectual. Some methods, such as the one presented in Section 5-1, would be insufficient by themselves to be used; however, the author includes the discussion to address as many potential options as possible.

Section 5-1: Individual Intrusion Detection System (IDS) Log Aggregation

Intrusion Detection Systems and their more advanced counterparts, Intrusion Prevention Systems (IPS), are used by system administrators across the country and around the globe. An IDS is a fundamental component to any network looking to have a modicum of security. The widespread proliferation of IDS's drives the effectiveness of this collection method. The method is almost trivial: aggregate IDS logs from individual systems and networks. A key component of such aggregation would be the simplifying and the consistency of form of data provided. IDS's log a sizeable amount of information about offending packets, and the aggregation of this data would be substantially easier if the data sent is limited to only that which is required by the database. Likewise, to facilitate the incorporation of new data a consistent form of reporting would reduce the processing time and difficulty. By providing a patch or application which automatically forwards the data to

the database, the developer could easy spread the request for information to IDS's worldwide with uniform, known reporting habits. This reporting methodology would mirror the approach used by the Internet Storm Center (ISC) as outlined in Section 2-3.

Two challenges are presented by this methodology: getting sufficient data and a selection bias. The first issue simply comes down to providing an incentive for local administrators to implement the data collection software. By making the content easily downloadable, convenient to install, and minimally invasive the designer can reduce the overall negative impact of the data collection process; however, the positive benefit to be gained from installing the patch is difficult to exemplify and thus encouraging administrators to implement collection requires using almost entirely normative, rather than legal, force. The second issue, a selection bias, results from the lack of complete implementation of the collection method and is the exact same issue the ISC suffers from. Since only some administrators will be implementing the data collection, more likely than not the administrators implementing the collection will be the ones suffering from foreign attacks, making the numbers received potentially inaccurate. Such a problem would make statistical correlation more apt to fall victim to error. For these reasons this method of IDS data aggregation is only recommended as used in conjunction with another method as simply aggregating disparate logs will fail to accurately depict the level and nature of cyber attacks.

Section 5-2: Industry Data Collection

Private companies have a strong focus on network security and monitor their data and servers carefully. In general, private industry has better logs of attempted intrusions into their private data than their civilian counterparts. This fact derives from their reliance on

private information for their profitability and future viability as a company. Private companies almost certainly use an IDS or IPS in securing their internal network and maintain logs for longer periods of time than common users.

Another method of collection would be similar to the method outlined in Section 5-1 but uses only private industry data rather than any administrator willing to implement the patch. Such a method would assume one of two scenarios: some sort of government requirement on private companies to yield their log data to the COCW database to ensure 100% compliance or surveying random selection of some subset of industrial companies representing a broad swath of fields to produce a statistically significant sampling group. In the first case, if such a requirement would be in place, the issue of gathering sufficient data would also be resolved as all data from private industry would be more than enough data to populate a database with information. In the second case, since the sample size has been determined to be statistically significant the data received would be sufficient to extrapolate as needed within reasonable limits.

However, two new problems arise with this methodology: imposition and privacy. The first issue for a developer would be the massive imposition on private industry. By demanding that such a monitoring and reporting software patch be applied to private industry in the first case the government would be imposing a financial and technological burden on private industry, a highly contentious point in itself. The second case would still require private industry to impose a company-wide patch or implementation; such an imposition would make many companies cringe and be extremely unlikely to participate. The second issue, that of privacy, will be discussed in depth in Section Six. Let it suffice to say that a company under attack would be less than willing to reveal such a fact to shareholders for fear

that such attacks may successfully injure the company (or that simply making the attacks public would worry potential investors). If such imposition and privacy concerns could be addressed, this method would be a quality method by which to collect the necessary relevant data without the inherent selection bias of the previous methodology becoming an issue.

Section 5-3: Internet Backbone Logging

The topology of the internet is not simply every computer connected directly to every other computer as one might imagine; rather connections form more of a tree diagram with long distance connections ultimately going through one of a small number of high speed nodes within the country. These high speed nodes compose the internet backbone, the routers which direct almost all traffic to where it needs to go *within* a given country. Any router on the internet backbone can pump an enormous amount of data, up to 4 gigabits per second (gb/s). These routers connect most systems together and handle a majority of internet traffic generally.

Applying a similar methodology as that from Section 5-1 would allow for the database to aggregate a large quantity of information relatively evenly. This method would essentially place an IDS on all internet backbone routers and aggregate the data into the COCW database. One major difference between the internet backbone and usual networks is that the internet backbone routers generally do not have an IDS set up. Therefore, implementing this methodology would require additional effort; most notably implementing an IDS solution on a router which has not previously had one. Once the IDS is set up, then the aggregation method outlined in Section 7-4 would be able to bring the data to the COCW database safely.

Two major issues arise with this methodology: inconvenience and data overload. Unlike the previous two implementation schemes, this methodology requires an IDS to be installed and worked into the already existing internet backbone framework. The inconvenience to installing an entirely new IDS is much greater than simply incorporating a software change on a previously integrated IDS. The private providers that host the internet backbone would most likely be less than thrilled at the prospect of being forced to integrate the new system. If, however, the installation was to go without incident and data was to be sent to the database in a uniform, continuous fashion (as seen in Section 7-4), then the database would have to consider the possibility of data overload. With the tremendous volume of attacks suspected to occur on a daily, even hourly, basis, the database would have to be able to handle millions if not billions of logs being sent each day. The processing power and memory space required to hold such data would be tremendous, requiring greater resources on the part of the database itself. One method of reducing the burden on the main database is through the use of substations as discussed in Section 7-4.

Section 5-4: Transnational Router Logging

Traffic on the internet comes from and goes all over the globe. These packets do not magically bridge the ocean; rather they use very specific high volume routing sites which transmit data between countries. Even data going from the United States to Canada goes through a limited number of specified routers.[26] This paper refers to these routers which

[26] While the exact number of routers remains unknown to the researcher, preliminary investigation suggests less than twenty transnational routers exist between the United States and Canada. This estimation was obtained by analyzing logs of transnational traffic between the United States and other countries gathered for the purpose of internet topology research. The logs were sponsored by The Cooperative Association for Internet Data Analysis, date to the year 2000, and can be found at <http://www.caida.org/tools/visualization/mapnet/Data/inc.txt>.

specifically send internet traffic between states as “transnational routers”. These transnational routers function similarly to the internet backbone in that they have a high level of traffic and do not currently have an IDS implemented generally speaking. The most notable difference between transnational routers and the internet backbone is the level of throughput. Transnational routers have potentially much higher speed for traffic, some exceeding 1 tb/s (roughly equivalent to 1,000 gb/s).

By tapping into these transnational routers instead of or in conjunction with the internet backbone, the database could more accurately gather the number of attacks going into or out of a country. This benefit derives from the nature of transnational routers: they are transnational. If an attack is detected coming into a French transnational router and that transnational router is only connected on that side to a Spanish transnational router, then one can conclusively say the attack came from the Spanish transnational router. This conclusion does not suggest the attack originally came from Spain, an important distinction; just that Spain was used in some way as a springboard for the attack. The largest advantage for this methodology (and the reason it is the most preferred option) is the fact that the use of transnational routers as monitoring points would provide for the minimum number of monitoring points to accurately determine the level of traffic entering and leaving a country. Transnational routers act as chokepoints for transnational internet traffic, thereby allowing for highly efficient monitoring.

The issues that arise in this method would be the same as Section 5-3, namely inconvenience and data overload. However, one more challenge would exist in this methodology: speed maintenance. The high level of traffic speed on these routers is important, and many IDS’s would slow traffic to a crawl or would not be able to keep up

with the speed of these transnational routers. As will be outlined in Section 7-2, though, proper implementation can prevent any sort of traffic speed reduction, removing this barrier to implementation.

Section 5-5: Route Recovery

Assuming that an attack was detected at some pair of points (such as two independent network IDS's found the same attack), the database could gather information as to the route taken to assist with the variable country-as-vector. Each router has a specific routing table it uses to translate a destination IP address into a destination hardware address. While these addresses change over time, if an attack is traced quickly enough (within 24 hours) by starting at the earliest known location of the attack the database could calculate the route taken by going router to router tracing the path taken. Such a methodology requires two known points of where the attack was. One point may be the targeted system if the attack was not stopped, allowing for any other point in the route taken by the attack to be the requisite second point of detection. This methodology would only be used in conjunction with another method, but may make a method such as the ones described in Sections 5-1 and 5-2 more detailed and useful for the database. This method would reduce the role of selection bias substantially as many more hosts not volunteering their information would have information volunteered for them thereby taking out the selection effect.

By using a route recovery algorithm the database will be better equipped to determine the common routes of attack (which networks the attacks go through) and thereby suggest which networks may be in need of stronger network defenses. Likewise, by seeing which networks failed to detect the attack the data collection mechanisms may be expanded when

the vector networks are informed of their need for increased security. The most notable downside to using route recovery is the sheer computational magnitude of the task. To check each router in the path, get its routing table, and continue to the next router takes quite a bit of time, time which the database would likely be unable to afford considering the streams of data that would be flowing in. The other major drawback is the selection bias may still exist in the route recovery method; if the two points of detection are always coming from the same two places then the routes recovered will more often be attacks against the same place, skewing data significantly. One method to offset the computational burden of route recovery would be using a separate server or database to compute route recoveries or to use some form of crowdsourcing, but such ideas exceed the scope of this COCW formulation.

Section 5-6: Recommended Data Collection Strategy

Though the method of implementation will ultimately be up to the developer, in theory the author recommends using a synthesis of the above ideas: the IDS technology described in Section 1-2 would be implemented on transnational routers and eventually filtered down to internet backbone routers on down to lower level networks. The reasoning for this choice is that transnational routers would provide the greatest level of relevant information, would have no selection bias, would be able to populate the database most accurately, and would best pinpoint the states from which attacks are coming. However, as the database sources proliferate and more countries start sending in data a need for solutions such as the one seen in Sections 5-1, 5-2, and 5-3 will arise from countries wishing to further pinpoint the source of cyber attacks within their own borders. This concept will be further discussed in Section 8-3. For the remainder of this paper, the author assumes that a

transnational router collection methodology is used for the sake of simplicity; however, such simplification need not rule out the use of alternate methods.

Section Six: Privacy

As with any data collection, there are major concerns regarding privacy that must be addressed before a database goes online. In the case of the COCW, some general privacy concerns include leaking information of who is bring targeted, censorship, and the preference for privacy of submission leading to database poisoning concerns. These issues will be looked at in the first section. Subsequent sections will outline three information categories which depict the different stages of data processing through the COCW submission process.

Section 6-1: General Privacy Concerns

Releasing the identity of who is actually being targeted by cyber attacks represents a tangible threat and a major privacy concern. Consider the scenario: the COCW declares attacks are taking place against a certain IP address. The owner of that IP address may be subject to major economic penalties if it is a public company; alternatively, the IP address may face increased cyber attacks as attackers bandwagon on the trend, further damaging whoever is at that IP address. The COCW absolutely cannot allow the identity of the target to be released. To combat this privacy concern, the author recommends that the information sent to the COCW from the substation (as outlined in Section 7-4) only include the country and industry data without any actual IP address being sent. By removing the destination IP information from the data sent to the COCW database then there is less of a threat of privacy being compromised. Sections 7-1 and 7-4 will outline the methods used by the COCW system to sanitize data and safely transport information to the database.

With any data monitoring system one must consider the threat of censorship or unlawful or unintended use of the system by the government or other body. In the case of the

COCW, the government could in theory use the monitoring infrastructure to prevent free speech by reporting the transmission of "suspicious packets" and changing the definition of "suspicious". This sort of change would be relatively straightforward and would require little time to implement once the COCW is receiving sufficient data. However, such a concern is outside of the scope of this formulation of the COCW database. The COCW uses no technology outside of the current purview of the United States government, and if the government wished to censor the populace in such a manner the challenge would be slight. The author contends that the lack of new or innovative technology being used suggests that, though censorship may be more of a threat in other states, in the United States such fears are only justified so far as they are justified absent the COCW system.

To ensure that those submitting information to the COCW are not threatened, one may quickly jump to using anonymous data submission similar to the system used by US-CERT; however, such a method would risk poisoning the database. Privacy of submissions is paramount to ensure data is submitted to the COCW. However, consider the following scenario: Country X sees the COCW labels it as having the most cyber attacks leaving its borders. To lose this poor reputation, Country X has one of its IDS's automatically submit thousands of false reports "detecting" attacks against Country X leaving Country Y. This information would poison the data in the COCW, rendering it useless or, worse, a tool of diplomatic deception. To prevent this possibility, the author recommends limiting the number of locations allowed to submit information. The author suggests a symmetric key encryption system to transmit the data would enable both the database to verify the submitted information while reducing the risk of retribution on the submitter due to the encrypted nature of the data being sent. The technical details of the use of symmetric key are outlined

in Section 7-3; however, for this section let it suffice to state that privacy of submission would be maintained while allowing for a receipt of submissions to be stored, thereby providing a disincentive for false reporting.

Though privacy is a great concern of the COCW, the author contends that through careful planning and the use of symmetric key encryption these concerns can be addressed constructively. Though the fear of censorship remains a constant one, the author can only recommend that the use of diplomatic channels would be the most likely candidate to prevent such a travesty from occurring.

Section 6-2: Information Breakdown

To understand the submission process of the COCW, one must consider the three major types of information in the COCW collection process. First, there is internal data, the data actually collected from the IDS. Second, there is the stored data, data kept on the substation or main database servers. Finally, there is the released data, the actual public data. These three types of information each have different levels of detail and thus require different amounts of protection from release. These differences will be outlined in the following three sections. For more details about where each type of data falls in the process, reference Section 7-4.

Section 6-3: Internal Information

The first type of information that needs to be considered is the information that will be collected by the IDS at the monitored point. This information is labeled “internal information” as it is data viewable only by the person who owns the network that is being monitored. In the case of a transnational router, for example, the private owner of that router

would be the only person with access to this internal information. Internal information requires the greatest protection against viewing as it contains all of the information available throughout the process, such as who is being attacked, where the attack is claiming to come from, and how often a given point is being targeted. This information, as discussed in Section 6-1, could be extremely damaging and must be protected at all costs. Therefore, internal information has the greatest need for protection, likely through strong encryption algorithms for any local storage. Then, to prevent this information from being released upon being transmitted, the data must go through a period of "data sanitization". This sanitizing process, outlined below in Section 7-1, will remove all identifying information from the internal information so when it is passed along nothing valuable would be gleaned from it by attackers. To further protect during transmission, the use of a symmetric encrypted channel, described below in Section 7-3, will prevent any intermediate users from reading the submitted information. In these ways the internal information will be protected to the greatest possible extent as loss of this data would be potentially catastrophic.

Section 6-4: Stored Information

"Stored information" is the data which has reached the substation or central database of the COCW (see Section 7-4 for substation information). This information is stored on the servers hosting the substations and central database and represents the sanitized information provided by the monitoring systems. This information is less secretive as all potentially harmful information has already been removed. However, given enough information attackers could engage in correlative reasoning to determine potentially harmful truths about the dataset. Therefore, though this information is not as detailed as internal information

stored information must also be protected from intrusion. Since there are less substations than deployed IDS's, the substations can afford greater protection measures and more vigilant monitoring by the COCW operators. Thankfully, the reduced value of the stored information makes it less likely to be targeted, but to ensure privacy protection will still be paramount in the design and operation of the COCW substation and central database servers.

Section 6-5: Released Information

The "released information" is the public-facing data available to anyone who desires access; the variables outlined in Section Four are examples of "released information". This information is especially important as it provides the crucial element of this database: free access by the public. However, this data is also the most vulnerable as "the public" can include attackers trying to get an edge or benefit from the data maliciously. This information, unfortunately, cannot be protected due to its public nature once it is released; therefore, the COCW must strive to make this information sanitized and safe *before* reaching this stage. To accomplish this objective, the use of data sanitization in the first two steps and the continued analysis of available information will be used to ensure that the released information does not yield data that ought to be kept hidden.

Section Seven: Data Collection Technical Details

The challenge of collecting data through the use of IDS logs is sizeable. In order to maintain the secure submission and collection of logs, a number of technical components are required. This section will begin by explaining how data will be sanitized to prevent dangerous data leakage. Then, the section will describe the many options available for high speed IDS implementations using FPGA's as presented in a variety of academic journals. This section will recommend a particular IDS implementation to use on connections between 1-10 gb/s. The use of such a high speed IDS will allow for the deep packet inspection required by the COCW database to be effective. Following this the author will discuss using symmetric key encryption to verify both the identity of the submitter and the privacy of submitted data and illustrate the benefits such a scheme would have. Finally, a depiction of the step-by-step data submission process will round out the section and explain each detail of the data submission process. Throughout this section the author recommends viewing Appendix B as an outline of the data submission process.

Section 7-1: Data Sanitization

To sanitize data, two possibilities need to be considered. The first possibility is that the network owners upon which the IDS's sit do not wish to release any destination IP information to the COCW system (a fair point, the issue is a major privacy concern as outlined in Section Six). In this case, the inference of the targeted country would be done at the IDS stage, requiring each IDS to take the destination IP data, lookup which country corresponds to that address, and then only send the country name on to the COCW substation. The same process would need to be done for target industry, attack

sophistication, country as vector, and any other inferred variables as discussed in Section Four. This sanitization method would prevent data leakage from the transmission, but would be more computationally intensive on the many endpoints of the system. The second possibility is if network owners trust the substation transmission security, then the COCW could do all IP lookups from the substations. This methodology would allow for less systems to spend time looking up corresponding countries for addresses and hopefully ensure the security of this lookup process. Either method would pass on the sanitized, non-descript country information to the next system in line. Appendix B depicts the first scenario, but the only change that would be made in the second scenario is that the IP addresses would be translated at the substation instead of the IDS. The remainder of the data (as outlined in Section 4-1), source hardware address, attack type, and attack name, would not require as intense of a sanitization due to its less harmful nature. For the remainder of this paper the situation shown in Appendix B as the author suspects the distrusting nature of the IDS owners will overrule the computational benefits of inferring the data at the substation level.

Section 7-2: High Speed IDS Recommendations

One major issue presented by data collection is the speed of traffic the IDS's would need to monitor. Transnational cables run at speeds up to or even exceeding 1 tb/s, a speed no packet searching IDS can handle. However, a solution does exist for this problem. Immediately upon reaching the destination country a high speed cable is split into separate lines which will then continue to split and separate into the lines used throughout the country. These individual lines usually do not flow at greater speeds than 3 gb/s. By monitoring these first or second order of separation lines one can conclusively state that any traffic exiting

through these lines came from the country on the other end of the transnational router. Due to asymmetric routing, an attack may go through one set of routers while coming back through a different set; however, with full implementation the attack will be detected coming back even if the attack goes through a completely different path.

While most lines do not exceed speeds of 3 gb/s, to be safe solutions will be presented which can handle packet inspection at throughputs of up to and including 10 gb/s. A number of IDS designs have been created to handle deep packet inspection. In 2005 Lin Tan and Timothy Sherwood outlined an architecture which allows for deep packet inspection at 10.07 gb/s and below.[27] Their design uses a custom made device with can maintain specific bounds on worst case performance, allowing for a high speed, concretely bounded IDS implementation.[28] The design presented by Tan and Sherwood is by no means unique; a number of such solutions have been produced by researchers in recent years.[29, 30] By implementing one of these high speed solutions on the required lines the COCW could be easily implemented on routers of any volume using the splitting property. By being able to handle deep packet inspection the use of high speed IDS's would allow for more effective monitoring and thereby better results and detection rates. The author leaves the detail of which particular high speed IDS to use up to the implementer.

---

[27] Tan, Lin and Timothy Sherman (2005). A High Throughput String Matching Architecture for Intrusion Detection and Prevention. < http://pages.cs.wisc.edu/~isca2005/papers/03A-01.PDF>.

[28] *Ibid.*

[29] Young H. Cho, Shiva Navab, and William H. Mangione-Smith (2002). Specialized hardware for deep network packet filtering. In 12th International Convergence on Field-Programmable Logic and Applications. Cho, Navab, and Mangione-Smith propose alternative solutions which also use FPGA's for speeds up to 2.8 gb/s, enough speed to handle the substantial majority of internet backbone lines.

[30] Ioannis Sourdis and Dionisios Pnevmatikatos (2004). Predecoded cams for efficient and high-speed nids pattern matching. In Proceedings of the Field-Programmable Custom Computing Machines. Sourdis and Pnevmatikatos present a solution which handles throughput up to 9 gb/s with alternate solutions which range from 2-9 gb/s throughput.

Section 7-3: Symmetric Key Encryption and Benefits

As mentioned in previous sections, there is a need to ensure that the data transmitted to the database is both unchanged in transit and from a proper source. To ensure both of these principles, the author recommends the implementation of a symmetric key encryption system. The symmetric key encryption system would be distributed physically to the COCW substations. Normally such a distribution system would be infeasible due to the large number of individuals requiring a key; however, in the case of the COCW database there will only be a limited number of cleared and reporting substations and data collection sites. Therefore, such an initial distribution is entirely acceptable if slightly time-intensive. The physical distribution provides the added benefit of allowing the COCW database personnel to ensure that the data collection systems (the IDS's) adhere to all COCW standards and have the proper updates, thereby preventing potential tampering by private parties. As such, the physical distribution of symmetric keys is highly recommended.

Following the initial distribution a traditionally secure symmetric encryption scheme will be used such as AES with a key of sufficient length. The use of symmetric system is ideal for a number of reasons. First, it ensures that data is coming from a legitimate user as only data encrypted with the expected key corresponding to the source IP address will be accepted as legitimate data. If the encrypted message does not decode with the proper key then the database will not enter the information into its stores, preventing an improper user from spoofing a reporting site. Also, symmetric encryption is computationally less intensive than asymmetric encryption, reducing overhead on reporting sites, the substations, and the central database. Since the data is encrypted the only manner by which data could be altered would be if an attacker broke the encryption. However, if an attack can break a symmetric

encryption cipher of the type being used in the COCW database system there are much larger problems than simply the potential poisoning of the COCW database. Finally, the since all data submissions would be tied to a particular IP address and key future auditing of data submission would be able to verify from where data was sent. By using this physically distributed key start and a symmetric encryption system the COCW database system can protect its data in transit, assure the source of its data, minimize computational burdens, and record data submission for future auditing.

Section 7-4: Data Submission Overview

To provide an understanding of the data submission process, this section will provide a step by step overview of how data will travel from the moment the attack is detected at a monitoring point to the time it is publicly displayed on the COCW database. Throughout this process references will be made to the graphical depiction in Appendix B.

The first step in the process is the detection of an attack by the IDS. At this point the IDS will naturally collect a lot of data, most importantly the data outlined in Section 4-1. This data will be collected using deep packet inspection and attack signature matching. This data will need to be very secure as it contains the most sensitive privacy concerns and is considered to be "internal data" (see Section 6-3). At this point a series of inferences will be done as outlined in Section Four for different variables, translating destination IP address into a country code, source IP into a source country code, labeling the detecting router as a vector or the destination, and other such inferences. The author will not go in depth in describing this process beyond noting the likely need for an outside database to be referenced and the need for this process to be secure to prevent sensitive data leakage. Once the inference

process is complete, the original and inferred data will be sanitized and sent on to the nearest COCW substation via a symmetrically encrypted channel as outlined in Sections 7-1 and 7-3.

The next step in the process involves moving the data to a COCW "substation". The substation is a data server down the line from the main COCW central database (as seen in Figure 2 in Appendix A). COCW substations provide a host of benefits to the data collection system. First, they reduce the computational burden on the central COCW server by spreading the collection over multiple points. Second, they prevent a single point of failure security hazard, allowing the possibility of the central database being compromised without losing everything. Third, substations reduce traffic throughout the internet by reducing the distance data needs to travel to be submitted. Fourth, the substations act as secure lockers, storing any information for as long as needed to allow for quick query and response for historic data. Finally, the substation method can add to the security of the COCW system by being the only recognized legitimate data submission authorities. By refusing any other entity from editing the COCW the substations can provide another safeguard against tampering. Substations may make the implementation of the COCW more costly and burdensome, but at the same time they will provide a host of benefits and in the view of the author are critical in the implementation of the COCW system.

Once the data reaches the COCW substation, the data is immediately stored in the substation data servers for future recall and must be protected as "stored data". This storage process is largely left up to the implementer, but should be in some way easily accessible by future query. The data received will be sanitized further and then await being sent to the main database. To reduce traffic, the substations will report to the COCW central database once, possibly twice a day with updates to statistics for their region. By only reporting once

or twice daily the substations will not substantially add to the traffic of their area. This data reporting will be done via an encrypted channel as described above.

Once the data reaches the COCW central server the information, received in a consistent format thanks to the use of substations, will be dynamically updated to the COCW database information. This information will then be made public through whatever medium the implementer sees fit. This will end the data submission process with the data having gone from a detected attack, through a substation, and finally into the central database.

Section Eight: Implications of the COCW

The successful complete implementation of the COCW database and its data collection method has two major implications for the international community: statistics and possibly peace. The first part, statistics, would clearly derive from researchers having access to numbers previously unavailable when researching the relationship between cyber attacks and other factors. Two theoretic examples of the use of such statistics are provided below in Sections 7-1 and 7-2. The possibility of peace derives from an inductive methodology which is rigorously described in Section 7-3, which, if accurate, could revolutionize internet security.

Section 8-1: International Research Design

This section of the paper will detail a hypothetical research design which incorporates the use of some variables from the COCW. The research design assumes the COCW has collected the necessary data to make statistically significant results. The purpose of having such a research design is to illustrate one of many relevant, interesting studies which could be undertaken by researchers if such a database were to be created.

The research question being asked is what relationship, if any, exists between the level of investment in cyber-related infrastructure and the level of outgoing cyber attacks? The hypothesis, posited by the author, suggests three distinct outcomes will arise. First, there will be a minimum amount of cyber infrastructure required to engage in cyber attacks readily, therefore below a certain infrastructure threshold outgoing and incoming cyber attacks will be relatively low as there exists insufficient infrastructure to have targets of value or engage in attacks. Second, there will be a middle level of infrastructure in which outgoing

attacks will increase due to the ability of attackers to engage in cyber attacks but incoming attacks will remain low due to the low value of domestic targets. Finally, the top tier of cyber infrastructure development will have a relatively equal amount of outgoing and incoming cyber attacks as the value of targets domestically will balance with international targets giving attackers equal incentive to attack internally as internationally. For a graphic depiction of this hypothesis, reference Figure 1 in the Appendix.

To test this hypothesis, a number of variables will be taken into account. The researcher will use the level of cyber infrastructure as the independent variable. The hypothesis is drawn from a theory that infrastructure levels affect cyber attack levels, so such a choice is logical. Cyber infrastructure will be measured by using the level of infrastructure investment in a country from data gathered at the International Telecommunications Union's International Communication Technology indicators database.[31] To measure the accuracy of the hypothesis, the researcher will look at the level of incoming and outgoing cyber attacks in the given state. This data will be drawn from the COCW, specifically the variables of country of origin and country targeted. A number of variables will be controlled for, including GDP, population, geographic size, and government type.

By using the data present in the COCW the researcher can investigate this obviously interesting claim. If the claim proves true, then there exists a rational self-interest for the most industrialized countries to pay less industrial countries to build their cyber infrastructure so criminals will target domestically rather than internationally thereby reducing the burden

[31] The International Telecommunications Union provides a host of measurements as to a country's level of investment in internet-related infrastructure. The author refers any interested parties to the ITU's website on international statistics at <http://www.itu.int/ITU-D/ict/statistics/>.

of attacks on the industrialized country. Such a conclusion would only be reached using the COCW database, illustrating the major benefit such a database would have in research.

Section 8-2: Target Research Design

Consider the question of who is under attack currently by cyber criminals. The research question being asked is if the new policy proposed by US Cyber Command moving military networks to a specific block of IP addresses has reduced targeting of non-military computers. The hypothesis presented by the author is that due to the increased accuracy of knowing where military computers are less non-military cyber attacks will occur but military cyber attacks will continue.

By using the COCW, one could determine what industries are being attacked by using the "industry targeted" variable. The COCW would provide this information for a long period of time both before and after the policy has been implemented. Using this data a researcher would be able to statistically test for a correlation before and after the policy implementation to see if the level of attacks have changed. If the change is not found, perhaps other changes have occurred to some variables; a researcher could easily test the other COCW variables to see what changes if any have been seen by the monitoring system following the policy being implemented. Such research would allow for cyber policies to be evaluated objectively and improved over time.

Consider the many possibilities from knowing which countries are the greatest targets internationally. Being able to unequivocally state that the United States is the number one target of cyber attacks in the world is a powerful statement. Alternatively, stating that countries such as Nigeria or Turkmenistan currently have almost no incoming cyber attacks

may compel attackers to reduce the burden of cyber attacks on the highest ranked country in order to pillage the relatively low traffic countries. Even the knowledge of seemingly irrelevant information such as which time of year has the highest number of cyber attacks would provide valuable research information for both the government and every day individuals hoping to maintain information security.

Section 8-3: Inductive Theory of COCW Proliferation

In this section the author posits a strong inductive case for how the public nature of the COCW in conjunction with the monitor's ability to discern at least one step of the path of the attack could lead to complete domestic implementation and possibly eventual complete international implementation of the COCW monitoring system.

Section 8-3-1: Assumptions

The opening assumptions of the inductive theory are relatively straightforward: first, the COCW has been implemented in some manner on a number of key routers throughout the United States (be these internet backbone routers or transnational routers); second, that the COCW monitoring system can discern at least one step back along the attack path from where an attack was detected (the "last step rule"). These two assumptions are justified above in the data collection schemes presented and through the logging of hardware source address information as explained in Section 4-1. One final assumption is made: there are cyber attacks originating in the United States to targets which require going through these monitored networks. This assumption basically assumes that Americans also engage in cyber attacks and use either the internet backbone or transnational routers to some extent, a justified assumption in the view of the author.

Section 8-3-2: Theories

The following theories derive from the above assumptions. Without loss of generality, assume an attack occurs from within the United States and is detected at router R. Due to the last step rule, the COCW can determine the previous step of the attack which will be within the United States. From this point, the prosecutor or law enforcement will have incentive to pressure the last step network to implement the COCW monitoring system. This theory logically comes from the desire of law enforcement to catch the criminal and if the COCW is deployed on the last step network then law enforcement will be one step closer to catching the perpetrator. As more and more attacks come out of the same malicious network, the level of pressure to implement the COCW monitoring system will increase accordingly until the previous step router is monitored. If the same malicious attacker is detected coming from the same source, the process of pressuring and implementation will continue down the attack path one step at a time until the attacks stop and can no longer be traced, the attacker moves to a new location, or the attacker is caught. Thus the inductive effect: the COCW monitoring system will move one step at a time down attack paths within the United States to enable prosecution. This induction follows from the fact that the same reasoning to move the COCW monitoring system to the next step down the attack path will apply to the newly monitored network.

The most likely reaction from cyber attackers will be constantly moving and shifting which networks they go through in order to circumvent the COCW monitoring system. However, such an action on the part of attackers will actually lead to faster implementation of the COCW system as more attack paths lead to more widespread pressure leads to more

widespread implementing of the COCW monitoring system. The key conclusion to be drawn is that the due to the public nature of the COCW monitoring system newfound pressure can be exerted on networks allowing attacks to go through them. This newfound pressure derives from the specificity of the number of attacks allowed and the clarity of which networks these attacks go through, both of which are provided only by the COCW database. Law enforcement currently is only able to say "we detected this one attack going through your router"; with the COCW, now law agents can authoritatively state "we detected *X* number of attacks going through your router in the last *Y* days". This authoritative numeric value will increase the efficacy of pressuring networks to implement information security policies. Thus, as attacks continue domestically the COCW will spread along all attack paths domestically until total implementation is complete.

Internationally the same result should occur: the COCW monitoring system will spread along attack paths. The rationale is the same as above. If an attack comes into the United States from Canada, the United States can exert newfound pressure on Canada to implement the COCW monitoring system on their transnational router to find the attacker responsible. As attacks continue the United States will be able to pressure Canada into implementing the monitoring system further into *their* domestic networks (and Canada should have no problem with this action since they value network security). Once more, the COCW system creeps along the attack paths of cyber attacks, becoming increasingly prevalent until cyber attacks are forced to go through COCW-monitored networks. This induction follows from the logic above. Ironically, cyber attackers movement will accomplish what diplomacy cannot: the spread of the COCW system.

These two explanations illustrate how the COCW monitoring system would be able to proliferate domestically and internationally in a relatively short amount of time with the proper interest and pressure applied to those reluctant to increase their own cyber security.

Section 8-3-3: Conclusions

If the COCW is implemented totally within the United States, there would be an undeniable benefit to our country: almost no cyber attacks for which there exist signatures would occur without more complex obfuscation. Consider the incentives for criminals in the situation. If a criminal engages in a cyber attack they will be immediately found, logged, and alerted to the authorities by the COCW database. There will be evidence of their actions linked to their hardware address which (ideally) is admissible in a court of law. Whatever benefit they might gain from engaging in such a cyber attack, the certainty of discovery would make cyber attackers within the borders of the United States hesitate. Though obfuscation may reduce the detection rate of attacks, the author accepts such limitations in current detection methods (as discussed further in Section 9-1).

The same effect would occur on attacks entering the United States from abroad. Cyber attackers would know the United States COCW database could track their attacks to its border and would have reduced incentive to attack the United States. Other countries would then become the target of cyber attacks and, coincidentally, be more in need of the COCW monitoring system to prevent attacks on themselves. Thus the proliferation of the COCW would reduce the number of incoming attacks into the United States and reduce cyber attacks originating in the United States over time.

The eventual international proliferation of the COCW monitoring system would lead to a much greater effect. Implementation would drastically reduce transnational cyber attacks as any attack could be traced to its country of origin, dramatically hindering the ability of cyber attackers to effectively engage in aggression over borders due to the possibility of being found increasing substantially. Also important to note, researchers would have infinitely better information and reams of data points to use to help study the field of cyber aggression internationally. The COCW database would be well-stocked with information available for international studies on cyber affairs which would be an end within itself worth working towards.

Through inductive logic the author firmly believes that the COCW monitoring system, if implemented on the edges of the United States cyber infrastructure, will quickly deploy within the domestic confines of the country until many simple cyber attacks which the COCW can monitor are effectively ended. Such a strong disincentive to attack cannot be overcome in the eyes of the author. Likewise, through induction the COCW system will spread along attack paths internationally. Skeptics may claim that there would be a lack of sufficient international pressure to force such a monitoring system. However, the author responds by suggesting that if a fraction of the pressure used to force intellectual property laws on other countries was to be used to propagate the COCW system then complete international implementation could be achieved within a very reasonable timeframe. The induction theory relies on pressure and accuracy; the former can only be accomplished with the help of the government, but the latter is ensured by the COCW database design itself.

Section Nine: Incomplete Implementation

The above analysis explains the many benefits of the COCW if it was to be implemented *completely* domestically or globally and gives a rationale for why and how such widespread proliferation should occur. However, the author recognizes the many challenges to such a complete and thorough implementation process and wishes to address some concerns regarding the difficulties of implementation, most notably the possibility of missing attacks and the difficult of getting all routers to sign on. Finally, the author will suggest some alternative methodologies which would use a partial implementation to provide very powerful data to be used in future research.

Section 9-1: Missing Attacks

The first and most obvious drawback to the COCW collection methodology is the possibility that some cyber attacks simply will not be detected by the deployed IDS. There exist a host of methods of obfuscating attacks to make them undetectable to a modern IDS; for example, an attacker could simply fragment their attack code to such an extent that each fragment by itself does not appear harmful but put together at the destination the attack would still take place. In another scenario, the attacker could use multiple layers of encryption to hide the contents of an attack packet from the IDS. Some forms of attack, such as the recent Stuxnet worm used against Iranian nuclear facilities, are so advanced and previously unseen that no method of obfuscation is even necessary; an IDS can only detect attacks it knows about, and some highly advanced or brand new attacks will inevitably slip through the detection system undetected.

Such a limitation does not in any way imply that the COCW would be thereby impotent. The COCW would undoubtedly be limited to collecting information about cyber attacks which are not obfuscated and are not so new so as to be beyond the purview of the implemented IDS's. However, from the outset the COCW defined "cyber attacks" as attacks which could and would be detected by an IDS; therefore the acceptance of these limitations is entirely justified under the COCW scheme. One should note that this drawback can and will be mitigated in part by the continuing evolution of IDS technology, but such evolution will never fully reach 100% of cyber attacks simply due to the simultaneous evolution of cyber attacker methodologies. As such, the COCW will be "incomplete" by its nature, but the author notes that a lack of completeness in this regard does not substantially hinder the COCW's objectives and can be regarded as unfortunate rather than damaging.

Section 9-2: Incomplete International Implementation

While the author hopes and dreams of a day in which the COCW monitoring system would be deployed internationally in all nations, one must recognize the probability that this ideal moment will never occur. There will always be countries which refuse to implement a monitoring system not controlled by their domestic government. If there is incomplete international implementation, are there still benefits to be had from the COCW database? The answer is unequivocally yes. Even without complete international implementation, the COCW will still provide data as to the number of cyber attacks entering and leaving the involved countries, which of these countries are most targeted, and where the attacks seem to be coming from. This data would still enable unique, beneficial analysis of attack sources and targets as outlined in Sections 8-1 and 8-2, though the country pool would be reduced.

Furthermore, the COCW would allow those countries that have implemented the monitoring system to levy more pressure on those which have failed to do so, increasing the push for internet security globally. Finally, the public information provided by the database would enable those countries which have enabled the monitoring to better secure their domestic networks as described in Section 8-3. Even without complete international cooperation, those that do participate fully in the COCW system will get the many benefits the COCW system has to offer.

Section 9-3: Incomplete Internal Implementation

The greatest threat to the effectiveness of the COCW methodology is an incomplete implementation of the monitoring system within a single country. That is to say, if a country does not implement the monitoring system on all routers of the same level within its sovereign realm (whether these are internet backbone routers, transnational routers, or others depending on the implementation scheme chosen), then the COCW will be unable to definitively state the data as outlined above. If only one router is left unmonitored, there is no way to prove that all cyber attacks are not going through that router and thus the data is corrupted.

The author has a suggestion how to measure the effectiveness of the COCW if implementation is incomplete, but before outlining this suggestion the author wishes to emphasize the importance or preference for complete internal implementation. With complete internal implementation the COCW system is able to adequately collect data so as to state information with a reasonable level of accuracy. Without such implementation a major beneficial feature of the database is lost. The author highly recommends that if the

system is to be implemented a complete internal implementation ought to be aimed for at all costs, even if such implementation requires removing other features from the database.

Section 9-3-1: Calculating Effectiveness of Incomplete Internal Implementation

The author proposes a calculation of the effectiveness of an incomplete internal implementation scheme assuming that some set of countries has implemented the COCW system on some subset of transnational routers**.** To calculate the effectiveness of incomplete internal implementation, the author will present a number of calculations based on an assumption of 70% implementation. This number is arbitrarily chosen and, as will be shown later, shifts above or below this level will have potentially exponential effects. When the author refers to "70% implementation" he suggests that 7 out of every 10 packets going into or out of a country pass through a COCW monitor on average. The question of the effectiveness of an incomplete internal implementation in part depends on the number of countries which have engaged in this incomplete implementation method. For the purposes of calculation the author assumes a reasonable number of countries have adapted the system and these countries are found more commonly in industrialized states due to their prevalence of being targeted by cyber attacks.

At this point the author assumes an attack is leaving a country (hereby called country A) and targeting a separate country (called Country B). Country B is assumed to be an industrialized country such as the United States and Country A is likely a less industrialized nation, such as Iran. To measure the "effectiveness" of the COCW system, the author asks what percentage of attacks that could be attack will be missed. If a low percentage of attacks are missed, then the COCW system is "effective" as almost all of the relevant target data is

being collected. If a large portion of attacks are being missed, however, then the system is "ineffective" as much data is being missed and is more likely to be skewed or inaccurate. The question of the probability of the attack being detected depends on the number of countries as vector used which have the monitoring system deployed. For *X* countries used as a vector and having the monitoring system, the general equation for the probability of an attack from Country A to Country B *not* being detected is listed below:

$$P(\text{attack NOT detected}) = ((0.3)^{2*X}) * 0.3$$

The first part of the equation, $((0.3)^{2*X})$, derives from the logic that an attack using a country as a vector must pass through two routers both of which have only a 30% chance of missing the attack. The first router is encountered upon entering the vector country and the second upon exiting the same country. Note that even if a country is adjacent to a non-implementing country this probability does not change. The last part of the equation, * 0.3, comes from the fact that Country B has the COCW system monitoring 70% of the traffic as well, allowing only a 30% chance that the attack passes through a non-monitored router in its final approach.

Some important facts can be drawn from this equation. The most startling fact is the exponential nature of the equation. Indeed, as the number of vector countries having the COCW monitoring system implemented increases the probability of detecting an attack increases by a power of two. Such effectiveness, even lacking a high number of countries, make the COCW system quite effectual. To illustrate this, consider Table 1 in the Appendix which outlines the probability of an attack not being detected based on the number of vector countries an attack must pass through. If the attack passes through four countries with the

COCW monitoring system the probability of an attack reaching its target without being detected is extremely low. To put the probability in perspective, if one million attacks were sent from Country A to Country B through four countries with the COCW system implemented, on average five attacks would be missed. Therefore, spreading a partial monitoring system through major threat vectors would quickly increase the probability of attack detection. Alternatively, if only adjacent countries had the COCW monitoring system online (so only one vector country, assuming all adjacent countries are innocent) the probability of attack detection is still over 97%. This would imply that if only the USA and each adjacent country applies the COCW a decent chance of detection would still exist. The above equation illustrates that a relatively small subset of countries would be required to make the COCW monitoring extremely effective.

There are some assumptions this model uses worth noting. First, the model assumes that attacks have an equal probability of going through any given router in a country. In reality, this fact is simply not true; Country A is surely more likely to go through a certain subset of routers to reach Country B. This reality both benefits and detracts from the above equation as outlined. By reducing the number of reasonable routers, if one knows Country A is a cyber threat then the probability of detecting an attack can be increased by the adjacent country by monitoring routers closer to Country A more heavily. However, if Country A is not suspected there is a risk that the above probability is too generous. Knowing ones enemies and placing the monitoring system accordingly would be greatly beneficial, but the author suggests that the above estimate is a worst case scenario estimate. A further assumption made is that level of traffic and level of malicious traffic have a direct correlation (since coverage was originally measured by traffic coverage). If this assumption is flawed

then the probability above may be off, but such new information would require a substantial theoretic reworking outside the scope of this paper. The final assumption that should be addressed is the 70% implementation assumption. As can be trivially proven, an increase or decrease in percentage of implementation will increase or decrease the probability of an attack being detected exponentially. As such, if Country B was able to encourage nearby countries to increase their implementation even slightly the benefit gained would be great, further encouraging the expansion of the COCW. The generalized formula for an implementation percentage *K* with *X* vector countries is outlined below:

$$P(\text{attack NOT detected}) = ((1 - K)^{2*X}) * (1 - K)$$

Section 9-3-2: Incomplete Internal Implementation Conclusions

Incomplete internal implementation is potentially the greatest pitfall to the COCW system. Without complete implementation a lot of attack source data is lost due to a lack of statistical significance. However, this incompleteness does not relegate the COCW to uselessness; to the contrary, reams of information regarding the number of attacks and the destination of such attacks can still be gathered without complete implementation. As outlined in Table 1 in the Appendix, the probability of detecting an attack is extremely high with only a few vector countries and a reasonable percentage of traffic being monitored. This fact implies that data regarding the destination IP address could be gathered, allowing the COCW to provide information regarding which industries are targeted, what attacks are most common, and what times are the most or least common for attacks to occur. Thus, even with incomplete internal implementation the COCW can yield vastly important target data and, with luck, lead to further internal implementation over time.

Section Ten: Summary and Future Research

To produce meaningful analyses, researchers require data. Empirical observation is the fundamental building block of science; yet to date no empirical database exists regarding the interesting and critical area of cyber attacks. To rectify this problem, the author has outlined a theoretical database called the Correlates of Cyber Warfare, or the COCW. This database would detect and report cyber attacks at a variety of levels. It would provide empirical data on an underreported sector of international relations: cyber security. The author has outlined what databases exist, the variables the COCW database would provide, the methods of data collection it could employ, addressed the concerns of privacy that may be voiced, brought up the technical details necessary, and highlighted the powerful implications such a system could have whether implemented completely or not. The Correlates of Cyber Warfare could provide researchers with new data to fuel grand new theories of cyber relations and curtail transnational cyber attacks. The greatest potential for future research is to implement the COCW system on a small scale. Beginning with a small network of a few computers, future researchers could create a COCW-like database system to test the working parts. The ultimate goal of future research ought to be perfecting and implementing the COCW on the scale for which it was designed: for nation-states. However, the author recognizes such a system may take time to develop properly, but hopes by outlining this design to begin the process of creating a database addressing the vastly underdeveloped field of cyber relations.

Appendix A

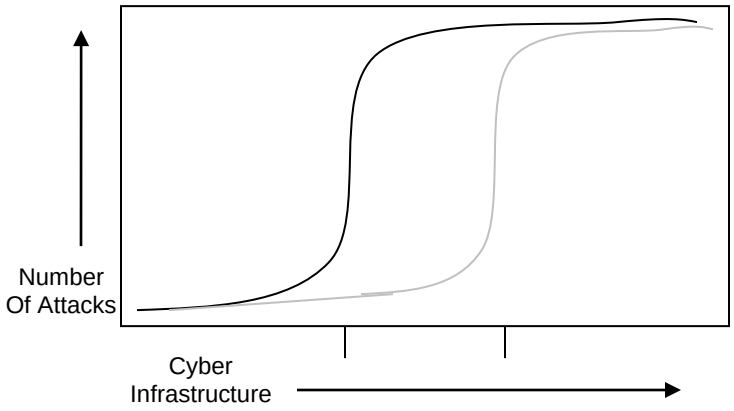

Figure 1: Hypothesized Relationship Between Cyber Infrastructure and Number of Cyber Attacks (Black = Outgoing Transnational Attacks, Grey = Domestic Attacks)

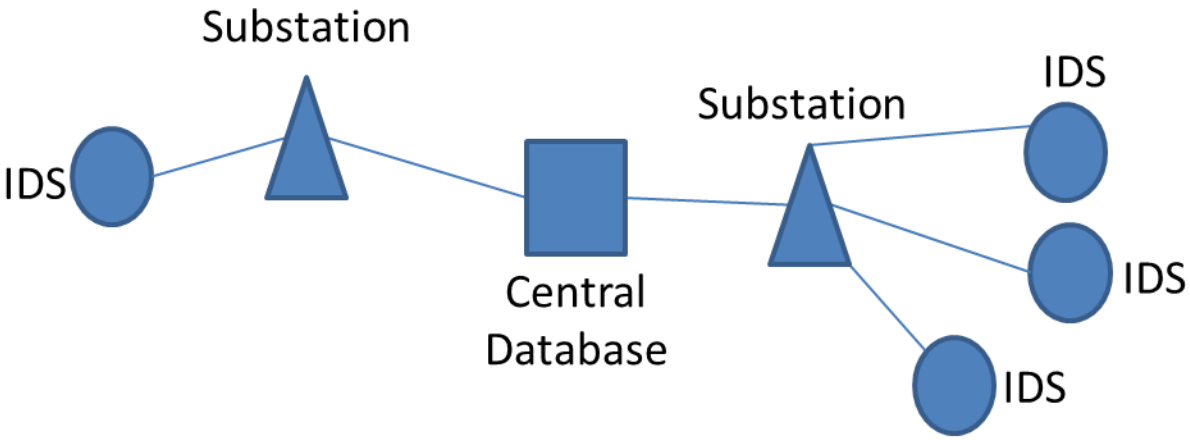

Figure 2: Correlates of Cyber Warfare Database Topology

# of Countries	Probability NOT Detected
0	0.3
1	0.027
2	0.00243
3	0.000219
4	0.00000590

Table 1: Probability of Attack Detection Based On Number of Vector Countries (Assumed 70% Implementation Rate)

Appendix B

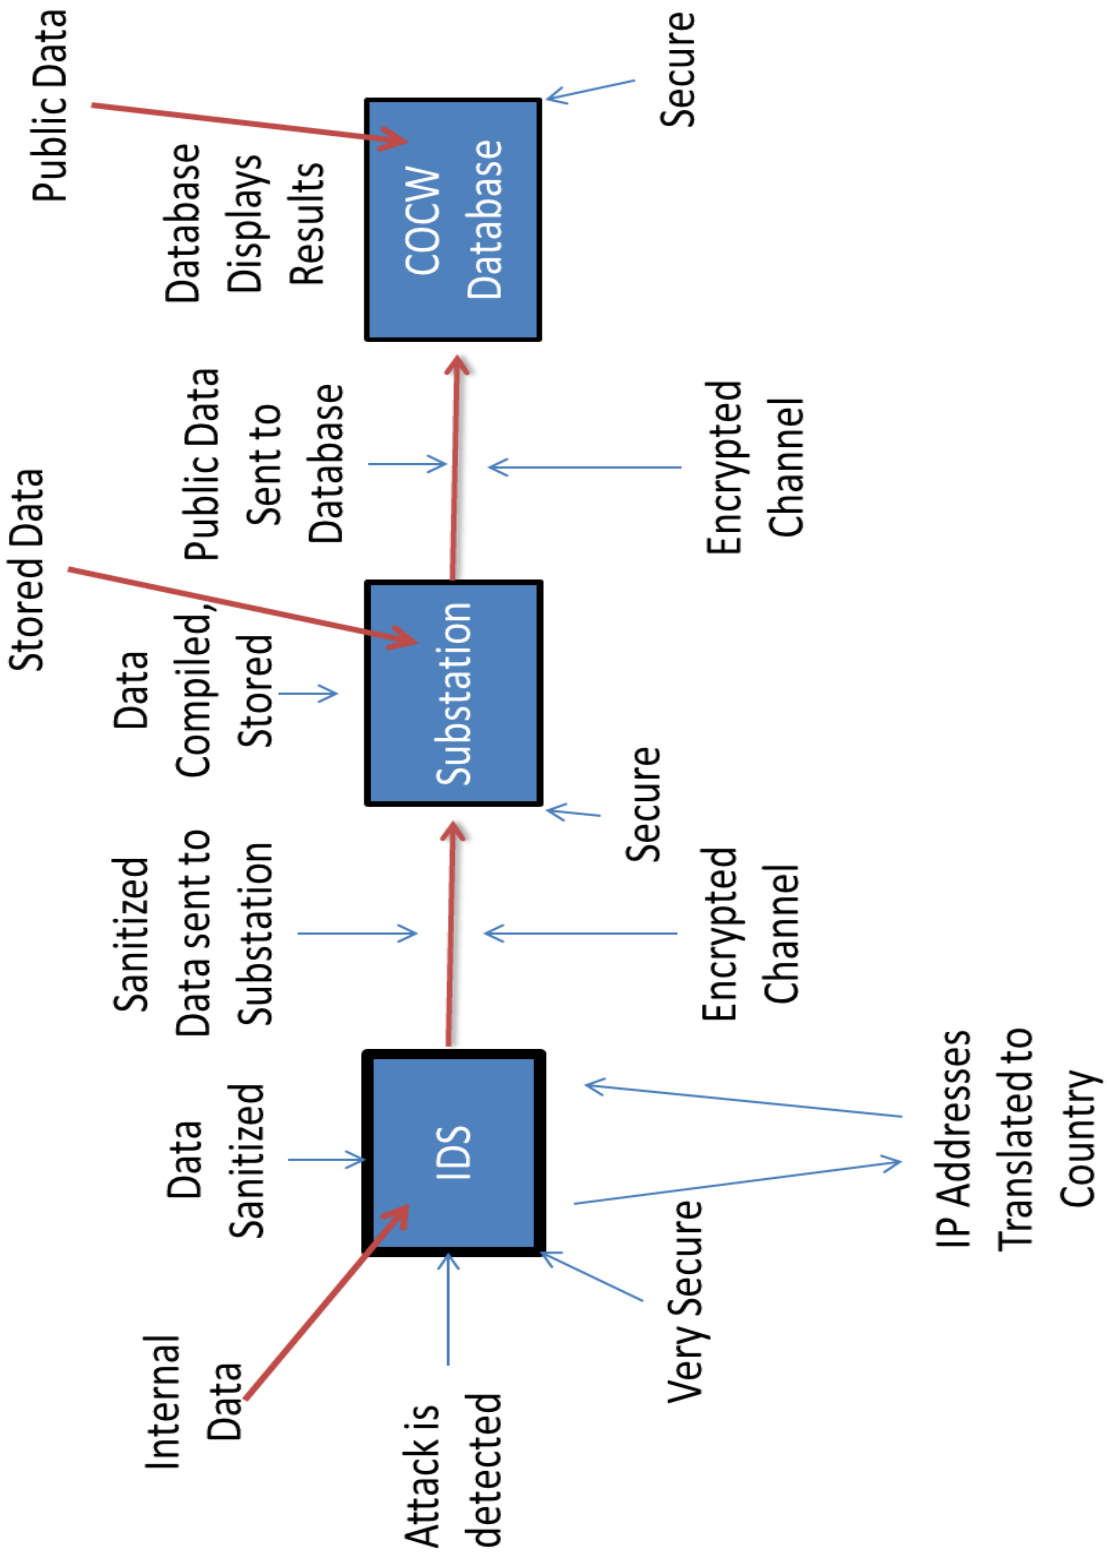
Public Data
Database Displays Results
COCW Database
Secure
Stored Data
Data Compiled, Stored
Substation
Public Data Sent to Database
Encrypted Channel
Sanitized Data sent to Substation
Secure
Encrypted Channel
Data Sanitized
IDS
IP Addresses Translated to Country
Internal Data
Attack is detected
Very Secure