Received 11 November 2022, accepted 2 December 2022, date of publication 7 December 2022,
date of current version 12 December 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3227420

APPLIED RESEARCH

A Zero-Day Cloud Timing Channel Attack

ROBERT FLOWERS, (Senior Member, IEEE)
Navy Federal Credit Union, Vienna, VA 22180, USA
e-mail: doc@stegascope.com

This work was supported in part by the College of Science and Technology, Bellevue University.

ABSTRACT The Intrusion Detection and Prevention System (IDPS) services of a North American cloud service provider were ineffective against a simulated network timing channel attack. During the tests, three conspiring white hat agents exchanged a total of 33,024 network packets. As the proxy based attack executed, the vendor's intrusion detection service did not generate a warning, nor did its intrusion prevention service drop packets. Throughout the experiment, 4,096 bytes of randomized data (simulating covert traffic) were exchanged over a 2.06 hour period (4.4 bits-per-second); however, the vendor's Artificial Intelligence (AI) enabled threat detection service did not issue an alert. A Wilcoxon Ranked Sum test on the before-and-after throughput confirmed none of the vendor's countermeasures triggered/intervened to a statistically significant degree (threat intel: $p = 0.703$, IDPS: $p = 0.998$, threat intel + IDPS: $p = 0.118$). These results indicate those accountable for data-oriented Service Organization Control (SOC) 2/3 reports (e.g., auditors, cybersecurity executives, etc.) should carefully examine the assurances offered by cloud service providers with regard to their network steganography defenses.

INDEX TERMS Network steganography, steganalysis, steganalyst, cloud, IaaS, timing channel, covert channel, data exfiltration, data theft, intrusion detection, intrusion prevention, IDPS, countermeasure, firewall, IP, TCP.

I. INTRODUCTION

More than a decade ago, Yale University researchers Ford and Aviram questioned the unmitigated trust corporations have of Cloud Service Providers (CSPs) [1]. As recently as 2018, Yale University researchers Deng et al. [2] warned the cybersecurity community about the threat posed by timing channels within the cloud. Those warnings echoed many of the concerns first expressed by their Yale University colleagues eight years earlier. A timing channel is a form of covert data transfer that uses time itself as a carrier and as a result, it is extremely difficult to develop, deploy, and detect [3], [4].

The Yale-identified threats distill into four inherent risks for cloud computing related to a timing channel: implicit clocks, shared resources, insider breach, and difficulty of detection. The last two of those inherent risks are the focus of the proxy based attack executed by this study. The National Institute of Standards and Technology (NIST) defines a zero-day weakness as a vulnerability within hardware/software that is discovered after its release [5], [6]. In the current context, the vulnerability in question is in the cloud. This study demonstrates timing channels in the cloud are, in fact, a zero-day security flaw.

As noted by MIT [7], organizations place insufficient importance on insider threats. Lack of emphasis on trusted bad actors compounds the issues arising from the implicit trust of CSPs. A system can perform well on any number of penetration tests, but its degree of porosity when insiders attempt to unlawfully move data from the inside out can be a cause for concern. When combined, the two threats create scenarios where an insider with high-level privileges, such as contractors and consultants, can steal confidential data and use the rapid setup and tear down of virtualization in the cloud to cover their tracks.

This research focused on the Infrastructure as a Service (IaaS) capabilities of a cloud service provider whose *infiltration* technologies were found to perform as advertised using assessment methods prescribed by the CSP.[1] Once those

The associate editor coordinating the review of this manuscript and approving it for publication was Abdel-Hamid Soliman.

[1] Future researchers can replicate the methods applied herein because the experiment did not rely upon any proprietary CSP features.

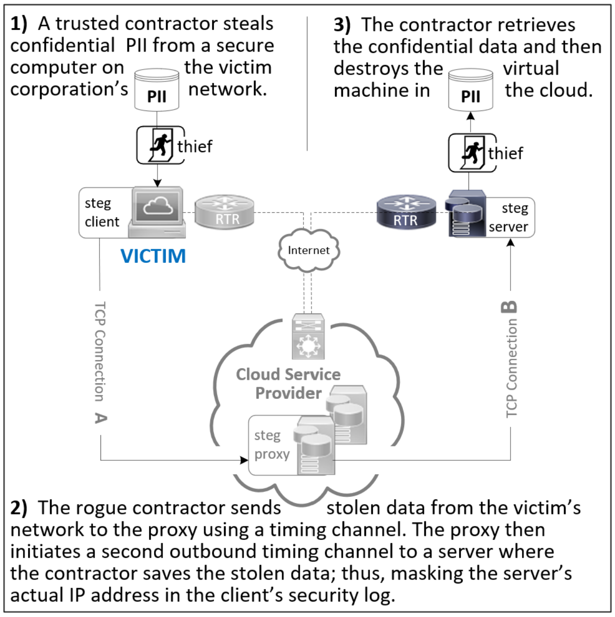
1) A trusted contractor steals confidential PII from a secure computer on the victim corporation's network.

3) The contractor retrieves the confidential data and then destroys the virtual machine in the cloud.

2) The rogue contractor sends stolen data from the victim's network to the proxy using a timing channel. The proxy then initiates a second outbound timing channel to a server where the contractor saves the stolen data; thus, masking the server's actual IP address in the client's security log.

FIGURE 1. **The *stegacloud* attack defeats client exfiltration defenses, thwarts the IDPS and threat intelligence controls in the cloud, and hides the true destination of the stolen Personally Identifiable Information (PII).**

infiltration countermeasures were confirmed, the study then interrogated those same defenses to determine how well they resisted *exfiltration* and alerted on attempts to hijack their IaaS layer by having them serve as a proxy during a sophisticated network steganographic assault. Figure 1 illustrates the vulnerability explored by this article.

Alerting is a foundational component of a timely security incident response [8]. Cloud deployments feature numerous security related advantages over typical on-premise solutions: especially those spanning multiple locations. One advantage is integration of event and incident notifications. Many enterprise-class CSPs offer Security Information and Event Management (SIEM) as an integral component to complement the inherent scalability and resilience benefits of cloud platforms. Whether part of the default package or a marketplace add-on, CSPs present homogeneous options for their clients to apply centralized management and monitoring functions to cover their entire end-to-end technical infrastructure.

Microsoft Azure Threat Intelligence [9], Amazon Web Services (AWS) Intelligent Threat Detection [10], Google Threat Intelligence [11], and similar CSP solutions are examples of smart countermeasures that dynamically adjust their operation based on external threat events. For example, assume Client A and Client B both utilize the same CSP. If an aggressor launches a Distributed Denial of Service (DDoS) strike against Client A, the CSP will not only execute countermeasures to defend against the originator of the Client A attack, it will also replicate the defensive configuration for Client B. Thanks to a multifaceted response, a single incursion could have the positive effect of hundreds or thousands of other customers becoming immune to the same offensive in a matter of seconds.

In order to determine the degree to which the intelligent responses described above could handle a real world data breach, the experiment that follows measured the CSP's response to suspicious Transmission Control Protocol/Internet Protocol (TCP/IP) segments when the selected CSP was not the direct subject of hostile action. The ubiquitous availability and rapid deployment of CSP virtual machines makes them an ideal target for indirect and ephemeral attacks via network steganography.[2]

II. BACKGROUND

A. STEGANOGRAPHY

Petitcolas [12] defined steganography as the practice of hiding data in plain sight. The word steganography (στεγανός, γράρ-ειν) has Greek origins. The root of the word is *steganos* which loosely translates to *cover*. The suffix *graphy* is based on the Greek word *graphein* which means *to write* [13]. Steganography, or hidden writing, is today a method of hiding communications so those looking at the output are unable to discern the cover (overt message) hides a second meaning (covert message) of which only the sender and receiver are aware. To an observer unaware of the covert content, the communication is interpreted based solely on the overt cover.

In the modern computing era, steganography is a tool with alleged links to global terrorism and documented evidence of international espionage. According to Schmurr and Crawley [13], Osama bin Ladin and his conspirators used steganography to plan the September 11, 2001 attacks as well as the bombing of embassies in Tanzania and Kenya. In sharp contrast to [13], Kellen [14] published a SANS Institute article indicating the direct links between bin Laden, the 911 attacks, and the application of steganography were tenuous. One crucial insight into the difference of opinion within the cybersecurity community on the terrorist-to-steganography connection was identified by the Dinca [15] survey. Dinca concluded the tools used by scientists to confirm the use of steganography by bin Laden were not sophisticated enough to detect the commercial-grade steganographic tools used by terrorists circa 2001.

Despite the varied perspectives on terrorist misuse of steganography, the nefarious application of steganography by Russian spies is undisputed. U.S. Department of Justice (DOJ) charging documents, such as United States v. Metsos, make it clear steganography has been a threat to U.S. national security for more than a decade [16]. For several years, the Federal Bureau of Investigation (FBI) conducted

[2] The experiment described herein was consistent with the principles outlined by Harvard University's Berkman Klein Center [18]. The CSP at the core of the current experiment invites researchers to perform testing of its detective/monitoring controls and does not require prior written notice of same. The 4.4 bit-per-second covert rate of each test was well within reasonable throughput thresholds to ensure there were no bandwidth impacts to co-located guests.

an investigation into spying by Russian Federation agents. The investigation concluded with the arrests of 10 spies who had been planted in various locations around the country including Arlington, Virginia and Seattle, Washington [17].

The FBI reported the goal of the Russian agents was "...to become sufficiently 'Americanized' such that they can gather information about the United States for Russia, and can successfully recruit sources who are in, or are able to infiltrate, United States policy-making circles" [19]. Section 3, subsection A.1.21 of the DOJ charging document describes how the Russian agents utilized Russian-created steganography tools to encode secret communications. The United States Government secured recordings of two of the defendants discussing how they used steganography to exchange classified data with their Moscow Center handlers.

B. NETWORK STEGANOGRAPHY

The deployment of steganographic tools by the Russian government makes it clear steganography is a vehicle for Advanced Persistent Threats (APTs). As a consequence, the level of sophistication of APT actors (e.g., hacktivists, federal governments, military forces, etc.) is far greater than the capability of the average hacker. Hosmer [20] stated highly skilled rogue agents have increased the complexity and detection resistance of the latest generation of data hiding methods. One of those methods is called *network steganography*. The SANS Institute [21] classified network steganography as one of the most complex forms of modern steganography [22]. Unlike the image-based steganography leveraged by the Russian Federation, network steganography uses network packets as cover data as opposed to the pixel values within image covers.

Basic network steganography techniques store covert data within unused locations inside network packet headers [23]. Such methods are called storage channels. As a more advanced form of network steganography, a timing channel does not physically store covert data, so investigators are unable to later perform forensic examinations to understand the nature of the breach [24]. A timing channel is considered the most complex network steganographic method because it adjusts packet transmission delays to hide data which increases the difficulty of both detection and prevention [25]. This state-of-the-art form of network steganography was the method applied by the experiment herein.

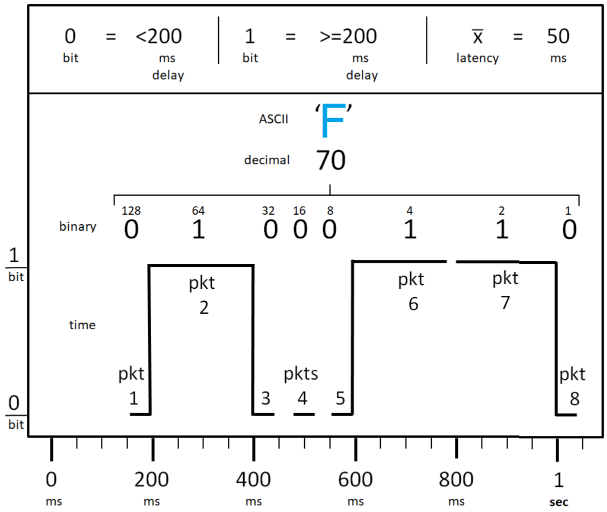

FIGURE 2. A timing channel places delays between network packet transmissions to encode data. Above, the sender transmits the letter 'F' by converting it to a decimal value of 70. The sender then transmits the binary representation of the number 70 one bit at a time in the form of delays between overt TCP segments. The receiver reverses the process. If the elapsed time since its previous reception of an overt TCP segment is less than 200 MS, the receiver stores a zero bit. If the delta is greater than or equal to 200 MS, the receiver stores a one bit. This process continues until the sender has transmitted all of the covert data.

C. CONVENTIONAL DATA REPRESENTATION

A computer represents data with binary code. Computer main memory, for example, can set an electrical charge above a predefined threshold voltage to represent the equivalent of a binary one bit. Conversely, if that charge is set below a specified threshold voltage or there is zero voltage, it represents a zero bit [26]. Four bits equal a nibble and eight bits is a byte. Those bytes can have a decimal value of zero through 255. The American Standard Code for Information Interchange (ASCII) uses the first 128 values of the 256 possible outcomes to represent the numeric values 0-9, the upper case letters A-Z, the lower case letters a-z, as well as punctuation and control characters [27]. Applying the ASCII 7-bit encoding scheme, the upper case letter 'A' is represented by the decimal value 65, the upper case letter 'B' has a decimal value of 66, and so on.[3] By employing recognized standards like ASCII and Unicode, each computer participating in a data exchange affords its user(s) a universally recognized mapping of computer values to human language elements. Without uniform standards, each computer manufacturer could have its own method for representing text that could be incompatible with others.

D. TIMING CHANNEL DATA REPRESENTATION

The mapping of data onto diverse mediums is not limited to computing [28]. There are many other ways to represent data. A common example is the oft-cited line from Longfellow's poem about Paul Revere: "One if by land, two if by sea" [29]. In the poem's encoding scheme, one lantern represented the enemy marching on land and two lanterns meant the British would invade via the Charles River. Per the Longfellow poem, light represented a storage medium for data. Time can also be utilized to transmit data [30]. As described by Keller [31] and Ganivev [32], a timing channel does just that by establishing temporal relationships between network packets so the receiver of those packets can determine whether it should decode a zero bit or a one bit. Figure 2 contains a line graph illustrating how time delays can be used to represent bits of data.

[3] Linux command shells and most text editors provide an example of how decimal-to-ASCII conversion works in practice. Hold down the Alt key while entering a two-digit decimal number between 65 and 90 using the numeric keypad on the *right side* of the keyboard. The equivalent ASCII character appears once the Alt key is released.

From the perspective of data thieves, the primary disadvantage of a timing channel is the level of effort required to create the algorithms needed to conduct a real-world exfiltration. A packet-level timing channel is the most difficult network steganographic method to code [33] because it requires a deep knowledge of the TCP/IP finite state machine during development as well as a thorough understanding of latency during execution [34], [35]. The programming languages used to create network steganographic programs also require considerable skill in order to generate IP datagrams and TCP segments [36], [37].

The payoff for a data thief willing to ignore the implementation difficulty is the degree of detection resistance demonstrated by a timing channel. Much like a sequence channel (i.e., a covert transfer mechanism that rearranges the transmission order of IP datagrams to encode covert data) the assailant need only alter the transmission pattern to create a timing channel. The result is a covert communication channel that is extremely hard to detect unless there are advanced controls in place to analyze statistical anomalies in the arrival times of network packets [38], [39]. Another advantage for the data thief, identified by Schmidbauer and Wendzel [40], is variations in packet timing introduced by multiple network hops can make detection even more difficult.

III. EXPERIMENTAL APPROACH

A. DESIGN

The current study posed the question: Are the most sophisticated cloud-based intrusion detection and prevention technologies capable of thwarting the most complicated data exfiltration attack? The null hypothesis (H_0) of the study stated there would be no difference in the successful transmission and reception of covert network packets regardless of the enabled or disabled state of cloud countermeasures. The research hypothesis (H_1) was the complement of the null hypothesis. The study applied a quantitative method with a before-and-after experimental design in order to determine the relationship between bits-per-second as the dependent variable and the state of defenses as the independent variable [41].

B. DEPENDENT VARIABLE

The bits-per-second (bps) dependent variable for each of the experiments was continuous. Pilot testing in preparation for the experiment showed experiments using throughput measurements, like bits-per-second, can generate negative or positive skew. Unlike prior tests, which were conducted on an isolated Local Area Network (LAN) by Flowers [42], the current study used multiple Internet hosts where latency, routing decisions, and the varying speeds of networks between client and server caused slight variations in data measurements [43]. The timing variations produced by routing decisions alone, as noted by Crepsi [44] and as previously mentioned by Schmidbauer and Wendzel [40], can have a normalizing impact on a time-based frequency distribution. Due to those variations, the skew in the current study was minimized but not entirely eliminated. To accommodate the residual non-normal distribution, an alternate statistical test was chosen to determine statistical significance.

C. INDEPENDENT VARIABLE

The independent variable selected was categorical. If the targeted CSP defense was disabled, the independent variable was equivalent to zero. If the targeted CSP countermeasure was enabled, the independent variable was recorded as one. The selected CSP's threat intelligence and IDPS services enabled the experimenter to drop suspicious packets, so the drop feature was enabled anytime the alert function was turned on. Enabling the dropping of suspicious network segments was necessary in order to ensure the dependent variable (bps) was impacted by the intervention of the CSP's countermeasures.

D. VARIABLE CONSIDERATIONS

A simple read-only alert may not have an affect on covert data throughput; however, it should be noted Rashid [45] found Deep Packet Inspection (DPI) deployed to detect network steganography can slow transfer speeds, so it is plausible the CSP's detective processing effort could have had an effect on the dependent variable. In either case, the selection of bits-per-second as a dependent variable was ideal because it can catch detective interventions as a side effect of slowed throughput due to DPI. Bits-per-second can also measure preventative intervention when dropped packets cause extended retransmission-based delays.

E. HIGH LEVEL ARCHITECTURE

Prior research on timing channel cloud vulnerabilities has focused on co-resident targets on the same cloud. As illustrated in Figure 1, the current research deviates from the requirement for co-residency and instead investigates the threat posed by a single cloud agent bookended by two non-resident client and server co-conspirators. Inserting indirection between client and server is not new. Indeed, using one or more intermediaries is the core identity obscuring benefit of The Onion Router (TOR) and therefore the dark web itself [46]. The value of indirection is such that the sponsors of TOR include the U.S. National Science Foundation and the U.S. Department of Defense (DOD) Defense Advanced Research Projects Agency (DARPA) [47], [48]. DARPA has funded the Advanced Research Projects Agency Network (ARPANET) which was the predecessor to the Internet, Messenger RNA research by Moderna which led to its COVID 19 vaccine, the Global Positioning System (GPS), and numerous other noteworthy technological innovations [49].

The indirection used within this study is more akin to the use of proxies to centralize all outbound web traffic by performing Source Network Address Translation (SNAT) on outbound IP datagrams [50]. Indirection, however, can have a darker side. The key difference between the application of a web proxy and stegacloud is the former is intended to

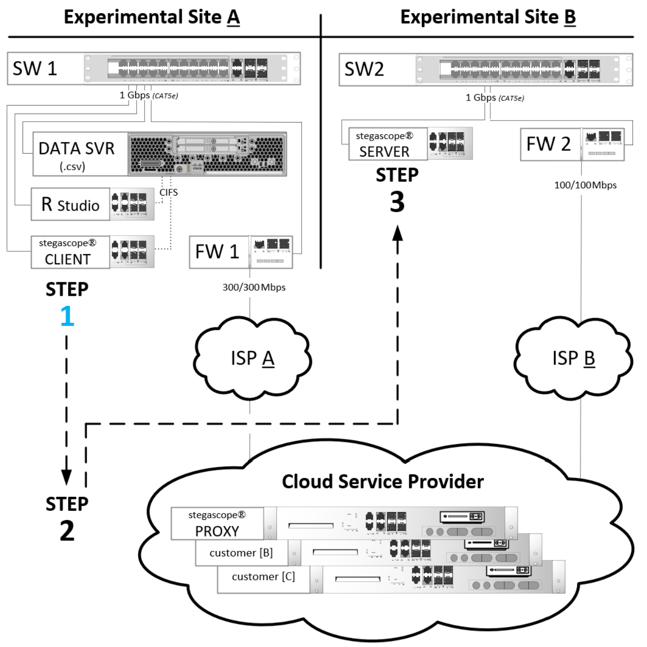

FIGURE 3. The experimental environment consisted of three core components: A client (Step 1) generating synthetic steganograms, an instance of the stegacloud proxy at the CSP's site (Step 2), and a server (Step 3) receiving the steganograms generated by the client.

preserve privacy whereas the latter is designed to simulate subverting it. The primary objective of the data thief is to obtain possession of confidential data. The secondary objective is to avoid detection. Unfortunately, stegacloud enables a rogue but trusted insider to accomplish both.

F. DETAILED ARCHITECTURE

As illustrated Figure 3, three distinct Internet locations were required to replicate a real-world environment. The first site hosted the client as well as test instrumentation. The second site was the CSP that housed the stegacloud proxy executable. The third and final Internet site consisted solely of the server component. Each site was protected by a firewall preventing extraneous computers from participating in the experiment. The active TCP ports were 8080 on the stegacloud proxy and port 80 on the host running the server.

During the experiments, the primary actors included the three conspiring agents described in Figure 1, but the number of participating TCP/IP hosts was expanded to properly measure the covert exchanges. A data server supplied storage services for the test executable operating on the client as well as the Comma Separated Values (CSV) files that held the outcomes from each test. Once all tests were complete, the R Studio statistical analysis application referenced those files via the Common Internet File System (CIFS). The CIFS server enabled the client and R Studio to access the test data via a network drive.

The server operated on a separate network connected to a second Internet Service Provider (ISP). The IP address of the server was in the Internet Assigned Numbers Authority (IANA) private range so the internal computers were not directly addressable from the Internet without traffic first going through a firewall [51]. The previously mentioned port 80 was opened on FW2. Destination Network Address Translation (DNAT) on FW2 enabled the stegacloud proxy to connect to FW2's public address, and port forwarding completed the configuration to facilitate proxy-to-server TCP segment exchanges. Similarly, the stegacloud proxy leveraged the CSP's DNAT functionality to translate client traffic received on its public IP address to a private IP address on port 8080.

TABLE 1. Statistical analysis indicates countermeasures were ineffective.

	Primary Test	Supporting Tests		Measurements	
Defense	$p - val$	Corr	Eff Sz	Mean bps	Samples
off				4.463	64
threat	0.703	-0.017	-0.033	4.453	64
IDPS	0.998	0.011	0.023	4.469	64
IDPS+	0.118	-0.129	-0.259	4.398	64

IV. RESULTS

A. STATISTICAL SIGNIFICANCE

As summarized in Table 1, a Wilcoxon Ranked Sum test of the CSP's threat intelligence service indicated no statistically significant difference ($p = 0.703$) between the inactive and active state of its defenses. An identical test demonstrated no statistically significant difference ($p = 0.998$) between the active or inactive state of the CSP's IDPS service. Similarly, a Wilcoxon test with both threat intelligence and IDPS enabled concurrently resulted in no statistically significant difference ($p = 0.118$) between the enabled and disabled state of the defenses. All tests were conducted using an alpha of 0.05 vs. 0.01 to minimize the risk of a false-negative; however, all quantitative output values were greater than the alpha, so the null hypothesis was accepted for each of the experiments.

Table 1 also summarizes the correlation and effect size tests. A Pearson Point Bi-Serial Correlation Test returned absolute values less than the range defined as 'no relationship' between the before-and-after state of all defenses [52]. The outcome suggests there was no evidence supporting a correlation between the categorical independent variable and the continuous dependent variable. The output of Cohen's D Effect Size tests corroborated the tests of statistical significance and the correlation tests. The effect sizes, for the individual before-and-after threat intelligence test as well as the individual IDPS service test, returned values in the 'no relationship' range. The concurrent threat intelligence and IDPS test produced a 'weak relationship' value. The mean of each group of tests was also measured based on the state identified in the *Defense* column in Table 1. The mean throughput values were measured with a high of 4.469 bps and a low of 4.398 bps.

B. TEST EXECUTION

During each test, the participating network hosts were monitored for anomalies that could cause overt TCP segment

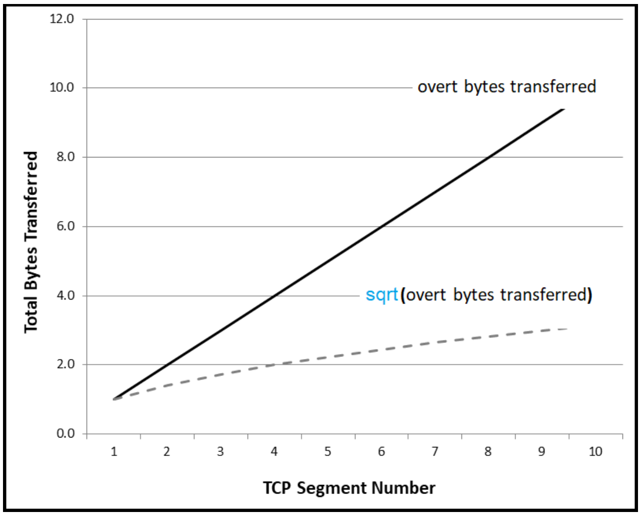

FIGURE 4. This line graph illustrates how packets can conform to the SRLS when sending each TCP segment, yet exceed the SRLS once a sample test is complete due to the sub-linear growth of square roots. Exceeding the cumulative SRLS can lead to false positives because the excess covert data increases detectability and gives an artificial advantage to the detective controls at the heart of the study.

delivery failure. The synthesized TCP stack utilized for each test supported all features of the TCP protocol necessary to deliver exceptional reliability including retransmissions, fast retransmissions, Selective Acknowledgements (SACKs), congestion control, and a sliding window tuned for high reliability as opposed to maximum throughput. The result was overt transfer effectiveness of 100% (with 100% accuracy) from client to server. The covert data transfer effectiveness was 100% (with 99.22% accuracy). The 0.78% covert inaccuracy was attributed to variations in routing decisions which negatively impacted TCP segment synchronization used by timing channels to covertly transfer data. The dependent variable did not factor for covert accuracy; however, 99.22% is a reasonable target given the inherent volatility of a timing channel. The following sections describe observations made during pilot runs of the experiment. Those observations led to adjustments to the testing procedures; however, the reported measurements were taken subsequent to those modifications.

1) SUB-LINEAR OBSERVATIONS

During testing, there were instances when the per-packet covert-to-overt ratio conformed to the Square Root Law of Steganography (SRLS), but the total number of covert bytes transferred exceeded the square root of the overt bytes transferred. The Square Root Law of Steganography (SRLS) states the number of covert bytes must be less than the square root of the number of overt bytes in order to lower the risk of detection [53]. The SRLS has implications for the researcher seeking to replicate real-world environments during controlled testing. Increasing covert payload size beyond the SRLS increases the chance a detective control will be triggered [54]. Conversely, the larger the covert payload, the greater the rate of covert transfer. Accordingly, there is a tension between covert throughput and covert detection risk.

TABLE 2. Covert transfer tests of threat intelligence and IDPS services conformed to the SRLS.

	Unit Counts		Max Covert	Conformance
Level	Overt	Covert	sqrt(Overt)	< SRLS
packet-level	8K bits	**1** bit	**89.4** bits	yes
sample-level	129K bytes	**16** bytes	**359.1** bytes	yes

2) SUB-LINEAR ROOT CAUSE

As shown in Figure 4, graphing the linear growth of overt values and their square roots was instructive. The graph illustrates the relationship between holistic totals (i.e., the cumulative number of overt bytes transferred at any point in the transmission process) and the square roots of those incremental values. The reason some tests demonstrated per-packet SRLS compliance despite holistic non-compliance was due to sub-linear growth. The count of overt bytes scales in a linear manner whereas the square root of the same value scales in a sub-linear fashion.

A review of steganography related scholarly literature revealed sub-linear growth of square roots was also analyzed in 2005 at the Cambridge University Computer Laboratory. Anderson [48] described sub-linear growth during an investigation into the theoretical limits of steganographic covers. Additional work by Filler et al. [55] four years later further explored the SRLS as it applied to digital image steganography. The relationship between holistic transfer totals and packet-level covert byte counts observed herein demonstrates prior research on detection resistance in digital image steganography provides valuable insight in network steganographic contexts.

3) SUB-LINEAR ADJUSTMENTS

To conform to the SRLS, the number of overt bytes was increased to ensure adequate cover data. Each TCP segment's overt data payload was increased to 1,000 bytes; thus, the covert-to-overt ratio for the timing channel in each test was one bit of covert data for every 8,000 bits of overt data: start-up packet notwithstanding. One thousand bytes of overt payload data was well within the TCP Maximum Segment Size (MSS) used for testing.[4] Equation 1 contains the SRLS constraint applied during testing.

$$c_{\text{covert units}} \ll \sqrt{o}_{\text{overt units}} \tag{1}$$

where c is the count of covert bits/bytes transferred from the sender to the receiver and *sqrt(o)* is the square root of the number of overt bits/bytes transmitted. Table 2 applies Equation 1 to the measured transfer totals.

An examination of Table 2 shows that at the packet level, the number of covert units (i.e., one covert bit per TCP segment) transferred was much less than the SRLS limit of

[4] MSS limits were conservative and below the maximum values specified during the TCP 3-way handshake to avoid ill effects if the CSP re-encapsulated data payloads. For example, the CSP's TCP stack could insert TCP Options between the TCP header and its subordinate payload. Doing so would have altered the segment count and inter-segment delays upon which a timing channel depends.

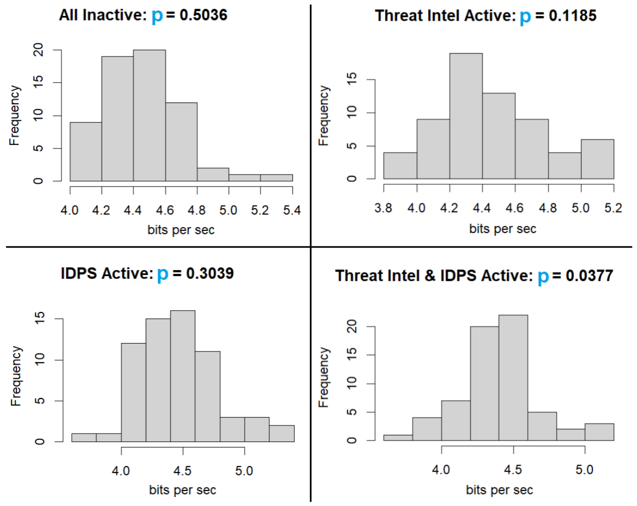

FIGURE 5. Despite apparent near-normal distribution in the visualizations above, the p-values reported by Shapiro-Wilk Normality Tests necessitated the use of non-parametric Wilcoxon Ranked Sum Tests.

89.4 bits. At the sample level, there were 64 samples for each of the four defense states listed in Table 1. For each of those samples, 16 bytes of covert data were transferred. Sixteen bytes of covert data is far less than the 359.1 byte maximum necessary to conform to the SRLS. Conformance to the SRLS constraints demonstrates no artificial advantage was given to the threat intelligence or IDPS countermeasures during the experiments.

C. TEST DATA

As shown in Figure 5, the data for some of the tests was normally distributed, but in other tests, data showed a small degree of skew. Shapiro-Wilk Normality Tests quantitatively confirmed the skew observed in the visualization of the data. Rather than perform transformations on the data to impart normality, a T-Test of Independent Means was replaced by a Wilcoxon Ranked Sum Test. Unfortunately, the Wilcoxon Ranked Sum Test is not as strong as the T-Test of Independent Means. To account for the difference in statistical strength, the ranked sum test was bolstered by the aforementioned correlation and effect size tests to ensure statistical significance was supported by complementary measurements.

V. DISCUSSION

In a stegacloud data exfiltration scenario, the victim organization controls just one third of the infrastructure components required for the attack. The rightful data owner does not control the data thief's stolen data cache, nor does the lawful owner of the data have direct control over the CSP. A proactive security organization does, however, have one critical indirect tool at their disposal to reduce the risks associated with the CSP's role in the attack described above: Service Organization Control (SOC) 2 reports. The following sections will describe how SOC reports can provide assurances the CSP has controls in place to mitigate a stegacloud attack.

A. SOC 2: OVERVIEW

As more on-premise data centers migrate to the cloud, instances of IaaS leakage will increase. Network steganography belongs to a large family of covert transfer techniques. The benefit of evaluating holistic data theft risk is no single technique will be overlooked. When enterprises evaluate prospective CSPs, they should ensure the service provider has an up-to-date SOC 2 report addressing a broad spectrum of covert threats. SOC 2 reports are issued by auditors who evaluate the design and operating effectiveness of controls within an organization promising safe and secure computing services [56]. An organization may have multiple SOC 2 reports if it sells numerous services. The SOC 2 report clearly defines the CSP's control objectives as well as any weaknesses for which the auditor issued a finding.

Gartner Research publishes guidance for organizations with regard to SOC 2 reports which is specifically applicable to those considering cloud services [57]. Based upon Gartner guidance, as well as the outcome of the network steganography tests above, consumers of cloud services should:

- Thoroughly review the control objectives to determine if data exfiltration controls (e.g., network steganography) are covered by the SOC 2 report.
- Check the User Control Considerations (UCCs) section of the SOC 2 report to see if the CSP places data exfiltration responsibilities on the client.
- Examine logging features and ensure virtualization deployment and destruction records are kept.

Unless an organization is an existing customer or is willing to sign a Non-Disclosure Agreement (NDA), obtaining a SOC 2 report may be difficult. Understandably, CSPs do not want to publicly broadcast security weaknesses. As an alternate first step, the American Institute of Certified Public Accountants (AICPA) suggests the use of a SOC 3 report [58]. It is a less private version of the SOC 2 report and is frequently publicly available. Ernst and Young [59] contains an example of a SOC 3 report. The SOC 3 report is a positive indication management's assertions have been tested by an independent third party. The AICPA [60] also issues comprehensive guidance on the UCCs described above. The key to these attestations and audits is that the CSP is evaluated against the full suite of AICPA Trust Services Criteria which consists of security, confidentiality, processing integrity, privacy, and availability [61]. The threat of network steganography intersects with several of those areas, so SOC 3 report language specifically addressing network steganography is a good sign the CSP has thought through all potential risks to the client's data.

B. SOC 2: ACCESS CONTROLS

All existing employees and contractors, CSP or otherwise, should have capabilities limited to the level of privilege necessary to perform their job function. Accordingly, the SOC 2 auditor will also inquire about the levels of privileged access held by current employees and whether the level of access is appropriate given the employee's job responsibilities. The 2022 Verizon Data Breach Investigations Report (DBIR) found 13% of breaches were caused by misconfigurations largely associated with improper cloud data access controls, so cloud access reviews (privileged or normal) are critical [62]. Certified evaluations by independent third parties ensures the risk-mitigating controls a CSP has in place are not simply one-time patches. SOC 2 reports give customers the confidence the CSP has had stringent processes in place for a period of time (i.e., the SOC 2 reports are retrospective, not prospective).

C. SOC 2: RECORDS RETENTION

The attack tested within this study focused on the misdeeds of a privileged user whose goal was to leverage the cloud as a stepping stone in a larger data exfiltration attack. Those criminal actions required the thief to become a CSP customer for just a few hours. The SOC 2 auditor will review the CSP's records of due diligence such as verifying the identity of its customers regardless of how long those customers actually subscribed. Another element of that due diligence is retention of all inbound and outbound TCP/IP connections made by/to a virtual machine running on the CSP's platform. Given the attack simulated within this study, law enforcement could utilize that connection history even after the attacker deletes the virtual machine to cover their tracks [63].

D. SOC 2: AUDITOR QUALITY

The prospective consumer of cloud services should also scrutinize the issuer (auditor) of the SOC 2 report. The Public Company Accounting Oversight Board (PCAOB) is the organization U.S. Congress empowers to ensure audit firms follow Auditing Standards (AS) when preparing their audit reports [64]. In 2015, the PCAOB issued an enforcement action against an auditor for failure to follow AS 7 (to be superseded by PCAOB AS 1220 in 2024 [65]) which requires an auditor to obtain an Engagement Quality Review (EQR) before delivering the outcome of an audit to a client [66], [67]. In this case, the client was a cloud-based provider of data storage services. The auditor was a sole proprietor and therefore likely unable to internally fulfil the requirement of having the partner-level secondary review mandated by AS 7. Size, credibility, and experience are essential attributes of an audit firm as it relates to assessing CSP countermeasure effectiveness given the technical complexity of network-based hidden channels.

VI. FUTURE RESEARCH

The architecture of the experiment conducted within the current study consisted of three primary TCP/IP hosts. Those hosts included a client, proxy, and server computers on separate networks. Future experimentation will investigate a two-tiered architecture with the client component, which represents the victim of the exfiltration, being located on the selected cloud platform. In such a configuration, the state of the exfiltration detection and prevention countermeasures on the client's cloud would serve as the independent variable. Another reasonable variant would be locating the server component on a cloud platform to examine the behavior of countermeasures when they are the destination of a timing channel exfiltration.

In either of the aforementioned test architectures, and unlike the current experiment, the CSP would serve as the direct object of the verification effort. Prospective consumers of cloud services could therefore leverage the analysis to improve their planned cloud deployments. Similarly, existing users of cloud platforms could leverage the experimental outcomes to bolster their current deployments via configuration changes. Such research could also inform their search for newly developed marketplace components specifically developed to mitigate the risk of timing channel exfiltration.

VII. CONCLUSION

The experiment conducted herein demonstrated the detective and preventative capabilities of the selected CSP are lacking with regard to network steganographic countermeasures. Ultimately, from a customer perspective, the client who is consuming the services of a CSP must ensure their end-to-end controls are sufficient to mitigate the risk of data exfiltration. Further, there is an intersection between public company cloud-based platforms and the United States Cyber Command's vision of achieving superiority in cyberspace. The stated objective is to match America's "…superiority in the air, land and space…" to its capabilities in the cyber domain [68].

The *Command Vision for U.S. Cyber Command* makes it clear adversaries of the United States of America plan to disrupt the economies of America and its allies as a means of warfare. Andress [69] noted the alarming fact that hackers benefit from the consolidation of disparate targets onto the cloud. As large enterprises continue to migrate currently dispersed operations onto the cloud-based resources of a relatively small number of large-scale cloud providers, an offensive against a single CSP could conceivably harm numerous unaffiliated corporations sharing the same cloud space. In addition to the two inherent risks evaluated by the experiment above, such centralization realizes a third Yale University identified inherent risk: shared resource vulnerability [1].

The current study revealed a considerable weakness as it relates to cloud platforms being used as an unwitting agent in a data theft. As noted by Yale University researchers, resource sharing and pay-as-you-go features are inherent to the cloud's core value, but those attributes are also prime vulnerabilities. Given the previously mentioned intersection between national and corporate weaknesses, it is in the common best interest of the private and military sectors to respond strategically to the threat of network steganography for the benefit of all.

ACKNOWLEDGMENT

The author would like to thank Dr. Mary Dobransky, Dean of the College of Science and Technology and Douglas Rausch, Director of the Center for Cybersecurity Education at Bellevue University for their continued support during the research, writing, and publication of this article. Additionally, the author expresses sincere appreciation for the support and guidance of Dr. Ian McAndrew, Dean of Doctoral Programs at Capitol Technology University.

REFERENCES

[1] A. Aviram, S. Hu, B. Ford, and R. Gummadi, "Determinating timing channels in compute clouds," in *Proc. ACM Workshop Cloud Comput. Secur. Workshop*, vol. 8. Chicago, IL, USA, Oct. 2010, pp. 103–108.

[2] (Jun. 2018). *Checking With Computation Tree Logic*. Hardware and Architectural Support for Security and Privacy. Los Angeles, CA, USA. [Online]. Available: https://caslab.csl.yale.edu/publications/deng2018cache.pdf

[3] F. Rezaei, M. Hempel, P. L. Shrestha, S. M. Rakshit, and H. Sharif, "Detecting covert timing channels using non-parametric statistical approaches," in *Proc. Int. Wireless Commun. Mobile Comput. Conf. (IWCMC)*, Aug. 2015, pp. 102–107, doi: 10.1109/IWCMC.2015.7289065.

[4] F. Benedetto, G. Giunta, A. Liguori, and A. Wacker, "A novel method for securing critical infrastructures by detecting hidden flows of data," in *Proc. IEEE Conf. Commun. Netw. Secur. (CNS)*, Sep. 2015, pp. 648–654, doi: 10.1109/CNS.2015.7346881.

[5] K. Dempsey, P. Eavy, N. Goren, and G. Moore. *Automation Support for Security Control Assessments*. National Institute of Standards and Technology. U.S. Department of Commerce. Accessed: Nov. 2, 2022. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/ir/2018/NIST.IR.8011-3.pdf

[6] M. Albanese, S. Jajodia, A. Singhal, and L. Wang. *An Efficient Approach to Assessing the Risk of Zero-Day Vulnerabilities*. National Institute of Standards and Technology. Office of Naval Research. Accessed: Nov. 2, 2022. [Online]. Available: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=913051

[7] N. Essilfie-Conduah, "A systems analysis of insider data exfiltration: A decentralized framework for disincentivizing and auditing data exfiltration," M.S. thesis, Syst. Des. Manag. Program, Massachusetts Inst. Technol., Cambridge, MA, USA, 2019.

[8] M. Cobb. *How Did They Get in? A Guide to Tracking Down the Source of APTs*. Accessed: Oct. 21, 2022. InformationWeek. [Online]. Available: http://twimgs.com/darkreading/advancedthreat/S4740412-howdidtheygetin.pdf

[9] A. McCollum, B. Gold, A. Buck, and Y. Levin. *Understand Threat Intelligence in Microsoft Sentinel*. Microsoft. Accessed: Oct. 21, 2022. [Online]. Available: https://learn.microsoft.com/en-us/azure/sentinel/understand-threat-intelligence

[10] Amazon. *Amazon GuardDuty Features*. Amazon Web Services. Accessed: Oct. 21, 2022. [Online]. Available: https://aws.amazon.com/guardduty/features/

[11] Google. *Configuring Threat Intelligence*. Google Cloud Armor. Accessed: Oct. 21, 2022. [Online]. Available: https://cloud.google.com/armor/docs/threat-intelligence#:~:text=GoogleCloudArmorThreatIntelligence,categoriesofthreatintelligencedata

[12] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Information hiding—A survey," *Proc. IEEE*, vol. 87, no. 7, pp. 1062–1078, Jul. 1999, doi: 10.1109/5.771065.

[13] A. Schmurr and W. Crawley, "Cybercrime in the United States criminal justice system: Cryptography and steganography as tools of terrorism," *J. Secur. Admin.*, vol. 26, no. 2, pp. 51–75, 2003.

[14] T. Kellen. (2022). *Hiding in Plain View: Could Steganography Be a Terrorist Tool?*. SANS Institute/GIAC. Accessed: 2005. [Online]. Available: https://www.giac.org/paper/gsec/1411/hiding-plain-view-steganography-terrorist-tool/102638

[15] L. M. Dinca, "Survey of the use of steganography over the internet," *Inf. Economica*, vol. 15, no. 2, pp. 153–164, 2011.

[16] (2010). *United States V. Metsos*. [Online]. Available: http://www.justice.gov/sites/default/files/opa/legacy/2010/06/28/062810complaint2.pdf

[17] (2010). *Ten Alleged Secret Agents Arrested in the United States*. [Online]. Available: http://www.justice.gov/opa/pr/ten-alleged-secret-agents-arrested-united-states

[18] Harvard University. *Google Cloud Penetration Testing: What It Is and How to Do It*. Berkman Klein Center for Internet and Society at Harvard University. Accessed: Oct. 21, 2022. [Online]. Available: https://cyber.harvard.edu/cyberlaw_winter10/Google_Cloud_Penetration_Testing:_What_It_Is_and_How_to_Do_it

[19] (2010). *United States V. Chapman and Semenko*. [Online]. Available: http://www.justice.gov/sites/default/files/opa/legacy/2010/06/28/_062810complaint1.pdf

[20] C. Hosmer. *Steganography: Chet Hosmer of Wetstone Technologies*. George Mason's Cybersecurity Innovation Forum. Accessed: Oct. 14, 2022. [Online]. Available: https://www.youtube.com/watch?v=sH3ZYx_WDMU

[21] E. Michaud. *Current Steganography Tools and Methods*. SANS Institute. Accessed: Oct. 26, 2022. [Online]. Available: https://www.giac.org/paper/gsec/2760/current-steganography-tools-methods/104695

[22] M. Raggo and C. Hosmer, *Data Hiding: Exposing Concealed Data in Multimedia, Operating Systems, Mobile Devices, and Network Protocols*. Waltham, MA, USA: Elsevier, 2013.

[23] T. Koziak, K. Wasielewska, and A. Janicki, "How to make an intrusion detection SystemAware of steganographic transmission," in *Proc. Eur. Interdiscipl. Cybersecur. Conf.*, Nov. 2021, pp. 77–82.

[24] D. Llamas, C. Allison, and A. H. D. Miller. *Covert Channels in Internet Protocols: A Survey*. Research Gate. Accessed: Nov. 2, 2022. [Online]. Available: https://www.researchgate.net/publication/245812185_Cover_Channels_in_Internet_Protocols_A_Survey

[25] H. Wang, G. Liu, and Y. Dai, "A detection method for cloak covert channel based on distribution of TCP burst size," *J. Inf. Hiding Multimedia Signal Process.*, vol. 6, no. 4, pp. 750–759, 2015.

[26] R. S. S. Schreibman and J. Unsworth, *A Companion to Digital Humanities* (Blackwell Companions to Literature and Culture). Hoboken, NJ, USA: Blackwell Publications, 2004.

[27] *How Unicode Relates to Prior Standards Such as ASCII and EBCDIC*. Oct. 20, 2022. [Online]. Available: https://www.ibm.com/docs/en/i/7.2?topic=wu-how-unicode-relates-prior-standards-such-as-ascii-ebcdic

[28] J. C. Wray, "An analysis of covert timing channels," in *Proc. IEEE Comput. Soc. Symp. Res. Secur. Privacy*, May 1991, pp. 2–7, doi: 10.1109/RISP.1991.130767.

[29] H. W. Longfellow. *Paul Revere's Ride*. Paul Revere Memorial Association. Accessed: Oct. 20, 2022. [Online]. Available: https://www.paulreverehouse.org/longfellowspoem/#

[30] G. Venkataramani, V. Prasadh, F. Yao, and H. Fant, "System and method for defense against cache timing channel attacks using cache management hardware," Patent U.S. 20 200 242 275 A1, 2020.

[31] J. Keller, "Multilevel network steganography in fountain codes," in *Proc. Eur. Interdiscipl. Cybersecur. Conf.*, Nov. 2021, pp. 72–76, doi: 10.1145/3487405.3487420.

[32] A. Ganivev, O. Mavlonov, B. Turdibekov, and M. Uzoqova, "Improving data hiding methods in network steganography based on packet header manipulation," in *Proc. Int. Conf. Inf. Sci. Commun. Technol. (ICISCT)*, Tashkent, Uzbekistan, Nov. 2021, pp. 1–15, doi: 10.1109/ICISCT52966.2021.9670109.

[33] T. Soni, "Moving target network steganography," M.S. thesis, Dept. Comput. Sci., Rowan Univ., Glassboro, NJ, USA, 2020. [Online]. Available: https://rdw.rowan.edu/cgi/viewcontent.cgi?article=3852&context=etd

[34] C. Kozierok, *The TCP/IP Guide: A Comprehensive, Illustrated Internet Protocols*. San Francisco, CA, USA: No Starch Press, 2005.

[35] N. B. Lucena, "Application-level protocol steganography," Ph.D. dissertation, Graduate School, Syracuse Univ., New York, NY, USA, 2009. [Online]. Available: https://graduateschool.syr.edu/

[36] B. Jankowski, W. Mazurczyk, and K. Szczypiorski, "PadSteg: Introducing inter-protocol steganography," *Telecommun. Syst.*, vol. 2011, pp. 1101–1111, Sep. 2011, doi: 10.1007/s11235-011-9616-z.

[37] W. Mazurczyk, P. Szaga, and K. Szczypiorski, "Using transcoding for hidden communication in IP telephony," *Multimedia Tools Appl.*, vol. 70, no. 3, pp. 2139–2165, 2014, doi: 10.1007/s11042-012-1224-8.

[38] J. Collins and S. Agaian, "Trends toward real-time network data steganography," *Int. J. Netw. Secur. Appl.*, vol. 8, no. 2, pp. 1–21, Mar. 2016, doi: 10.5121/ijnsa.2016.8201.

[39] H. Nafea, K. Kifayat, Q. Shi, K. N. Qureshi, and B. Askwith, ''Efficient non-linear covert channel detection in TCP data streams,'' *IEEE Access*, vol. 8, pp. 1680–1690, 2020, doi: 10.1109/ACCESS.2019.2961609.

[40] T. Schmidbauer and S. Wendzel, ''Hunting shadows: Towards packet runtime-based detection of computational intensive reversible covert channels,'' in *Proc. 16th Int. Conf. Availability, Rel. Secur.*, Vienna, Austria, Aug. 2021, p. 71, doi: 10.1145/3465481.3470085.

[41] C. M. Patino and J. C. Ferreira, ''Internal and external validity: Can you apply research study results to your patients?'' *J. Brasilieiro Pneumologia*, vol. 44, no. 3, p. 183, May 2018, doi: 10.1590/S1806-37562018000000164.

[42] R. Flowers, ''Performance impact of header-based network steganographic countermeasures,'' *IEEE Access*, vol. 10, pp. 92446–92453, 2022, doi: 10.1109/ACCESS.2022.3202556.

[43] A. Mileva and B. Panajotov, ''Covert channels in TCP/IP protocol stack—Extended version,'' *Open Comput. Sci.*, vol. 4, no. 2, pp. 45–66, Jan. 2014, doi: 10.2478/s13537-014-0205-6.

[44] V. Crespi, G. Cybenko, and A. Giani, ''Engineering statistical behaviors for attacking and defending covert channels,'' *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 1, pp. 124–136, Feb. 2013, doi: 10.1109/JSTSP.2012.2237378.

[45] A. Rashid, R. Ramdhany, M. Edwards, S. M. Kibirige, D. Hutchison, and R. Chitchyan. *Detecting and Preventing Data Exfiltration*. Lancaster University. Accessed: Oct. 22, 2022. [Online]. Available: https://www.researchgate.net/publication/299666817_Detecting_and_Preventing_Data_Exfiltration_Executive_Summary

[46] TOR. *Sponsors*. Accessed: Oct. 22, 2022. [Online]. Available: https://www.torproject.org/about/history

[47] DARPA. *About DARPA*. Accessed: Oct. 22, 2022. [Online]. Available: https://www.darpa.mil/about-us/about-darpa

[48] R. N. Anderson. *DARPA and the Internet Revolution*. Columbia University. Accessed: Oct. 22, 2022. [Online]. Available: https://www.researchgate.net/publication/330638395_DARPA_and_the_Internet_Revolution

[49] Economist. *A Growing Number of Governments Hope to Clone America's DARPA*. The Economist Newspaper Limited. Accessed: Oct. 22, 2022. [Online]. Available: https://www.economist.com/science-and-technology/2021/06/03/a-growing-number-of-governments-hope-to-clone-americas-darpa

[50] K. R. Fall and W. R. Stevens, *TCP/IP Illustrated: The Protocols*, 2nd ed. Upper Saddle River, NJ, USA: Pearson, 2012.

[51] R. Burk, M. Bligh, and T. Lee, *TCP/IP Blueprints*. Indianapolis, IN, USA: Sams Publishing, 1997.

[52] R. G. V. D. Berg. *Cohen's D—Effect Size for T-Test*. Accessed: Oct. 16, 2022. [Online]. Available: https://www.spss-tutorials.com/cohens-d/

[53] H. Y. El-Arsh, A. Abdelaziz, A. Elliethy, and H. A. Aly, ''Information-theoretic limits for steganography in multimedia,'' 2021, *arXiv:2111.04960*.

[54] B. A. Bash, D. Goeckel, and D. Towsley, ''Square root law for communication with low probability of detection on AWGN channels,'' in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2012, pp. 448–452, doi: 10.1109/ISIT.2012.6284228.

[55] T. Filler, A. D. Ker, and J. Fidrich. *The Square Root Law of Steganographic Capacity for Markov Covers*. Oxford University Computing Laboratory. Accessed: Nov. 5, 2022. [Online]. Available: http://www.ws.binghamton.edu/fridrich/research/Fil09spie.pdf

[56] A. Yawn. *An Expert's Guide to Reviewing SOC 2 Reports*. SANS Institute. Accessed: Oct. 13, 2022. [Online]. Available: https://www.sans.org/blog/expert-guide-reviewing-soc2-reports

[57] L. Leong, G. Petri, B. Gill, and M. Dorosh. *Magic Quadrant for Cloud Infrastructure as a Service, Worldwide*. Gartner. Accessed: Oct. 12, 2022. [Online]. Available: http://www.blackfinsquare.com/wp-content/uploads/2017/08/2016-Gartner-Quadrant-IaaS.pdf

[58] AICPA. *SOC 3: SOC for Service Organizations: Trust Services Criteria for General Use Report*. American Institute of Certified Public Accountants. Accessed: Oct. 16, 2022. [Online]. Available: https://us.aicpa.org/interestareas/frc/assuranceadvisoryservices/aicpasoc3report

[59] EY. *System and Organization Controls 3 (SOC 3) Report: Report on the Amazon Web Services System Relevant to Security, Availability, and Confidentiality*. Ernst & Young. Accessed: Oct. 16, 2022. [Online]. Available: https://d1.awsstatic.com/whitepapers/compliance/AWS_SOC3.pdf

[60] A. Katcher and R. Brown. *Understanding How Users Would Make Use of a SOC 2 Report*. Trust/Data Integrity Task Force. Accessed: Oct. 16, 2022. [Online]. Available: https://www.aicpa.org/content/dam/aicpa/interestareas/frc/assuranceadvisoryservices/downloadabledocuments/soc2-user-document.doc

[61] AICPA. *2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy*. American Institute of Certified Public Accountants. Accessed: Oct. 16, 2022. [Online]. Available: https://us.aicpa.org/content/dam/aicpa/interestareas/frc/assuranceadvisoryservices/downloadabledocuments/trust-services-criteria.pdf

[62] G. Bassett, C. Hylender, P. Langlois, A. Pinto, and S. Widup. *Verizon 2022 Data Breach Investigation Report (DBIR)*. Accessed: Oct. 16, 2022. [Online]. Available: https://www.verizon.com/business/resources/reports/2022/dbir/2022-dbir-data-breach-investigations-report.pdf

[63] R. Hunt, ''New developments in network forensics—Tools and techniques,'' in *Proc. 18th IEEE Int. Conf. Netw. (ICON)*, Singapore, Dec. 2012, p. 381, doi: 10.1109/ICON.2012.6506587.

[64] PCAOB. *About*. Public Company Accounting Oversight Board. Accessed: Oct. 24, 2022. [Online]. Available: https://pcaobus.org/about

[65] PCAOB. *AS 1220: Engagement Quality Review*. Accessed: Oct. 24, 2022. [Online]. Available: https://pcaobus.org/oversight/standards/auditing-standards/details/AS1220

[66] P. W. Brown. (2015). *Order Instituting Disciplinary Proceedings, Making Findings, and Imposing Sanctions*. Public Company Accounting Oversight Board. Accessed: Oct. 24, 2022. [Online]. Available: https://pcaob-assets.azureedge.net/pcaob-dev/docs/default-source/enforcement/decisions/documents/coons.pdf?sfvrsn=8fb4d5c5_0

[67] PCAOB. *Auditing Standard, no. 7: Engagement Quality Review*. Accessed: Oct. 26, 2022. [Online]. Available: https://pcaobus.org/oversight/standards/archived-standards/pre-reorganized-auditing-standards-interpretations/details/auditing-standard-no-7_1836

[68] CYBERCOM. *Achieve and Maintain Cyberspace Superiority: Command Vision for U.S. Cyber Command*. U.S. Cyber Command. Accessed: Oct. 22, 2022. [Online]. Available: https://www.cybercom.mil/Portals/56/Documents/USCYBERCOM VisionApril2018.pdf

[69] J. Andress and S. Winterfeld, *Cyber Warfare: Techniques, Tactics and Tools for Security Practitioners*. Oxford, U.K.: Syngress Publishing, 2011.

ROBERT FLOWERS (Senior Member, IEEE) received undergraduate and graduate degrees in information technology from American Military University, WV, USA, and the Doctor of Science (Sc.D.) degree from Capitol Technology University, MD, USA. He is an Adjunct Professor with Bellevue University, where he teaches graduate cybersecurity courses. He is currently an Assistant Vice President with the Navy Federal Credit Union, VA, USA, where he has worked for the last 25 years.