

1. Introduction

The Internet is a low-cost, low-risk, and high-value of return intelligence gathering and archival system. The Internet has been proven to have large amounts of information that is typically unintended for the intruder, but is freely available and relatively easy to discover.

Often users of computer systems perform operations on the Internet, such as searching for information, which they would like to keep anonymous. As a result, organizations (and users) have begun deploying ANSs. These include both software products such as Anonymizer [1] and open source projects such as Tor [2].

The focus of this paper will be on a variation of the ANS offered by Anonymizer, Inc [1]. These ANSs can be exploited by a variety of attacks that target applications as well as attacks that target the users.

This paper considers three main methods that could be used by an intruder to possibly discern the true identity of the user. The first method of potentially de-anonymizing a user is by a user clicking on a file that has logging capabilities enabled. This would allow the intruder to perform traffic analysis and find a correlation between the anonymized IP traffic and the user's real IP traffic. Next, the intruder could create a malicious file that when opened with a vulnerable application on a vulnerable operating system could exploit the user giving the intruder the ability to take control of the system. The last method is attributable searching while on a non-attributable ANS. By definition, the action of performing attributable search queries could yield information that might potentially de-anonymize the user. The intruder could gather attributable information based on search queries including but not limited to local weather, sports, and restaurants. Figure 1 shows the different de-anonymization methodologies discussed in this paper.

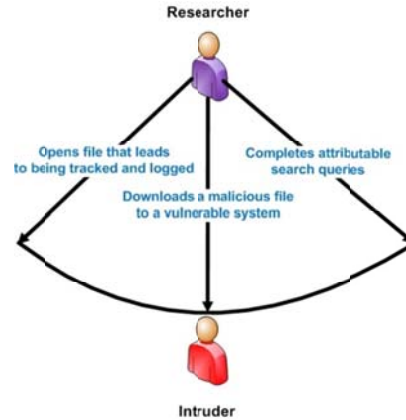


Figure 1: Possible de-anonymization methods

To gather experimental evidence we developed a case-study based on a typical ANS deployment in a large US-based organization. We call this version the case-study anonymizing network system or CSANS. CSANS represents a typical implementation of an ANS that is available today.

This experiment required the research manager of the case study organization to design a fictitious research task. The research task was designed entirely by the research manager with no assistance from the authors of this paper. Therefore, we conducted the experiment on a real operational environment and gathered the reported results from field experiments using real human subjects.

Figure 2 shows the relationships between the key players and the entities associated with the experiment.

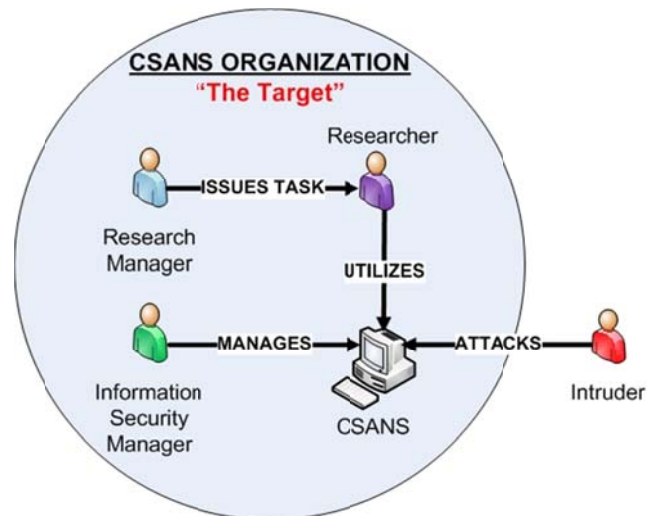


Figure 2: Relationships between key players in the experiment

As depicted in Figure 2, the research manager issues a specific task to the researcher which will require the use of the CSANS to accomplish. The information security manager is responsible for developing, securing, monitoring, and overall managing the CSANS. The intruder is assumed to be someone outside the organization who is interested in breaking the anonymization of the CSANS organization and its users in an effort to gain the researcher’s personal identifiable information (PII). While it is possible that the intruder is an inside threat we assume that the intruder is not associated with the CSANS organization. Figure 3 shows the interactions between the key players and the entities associated with the experiment.

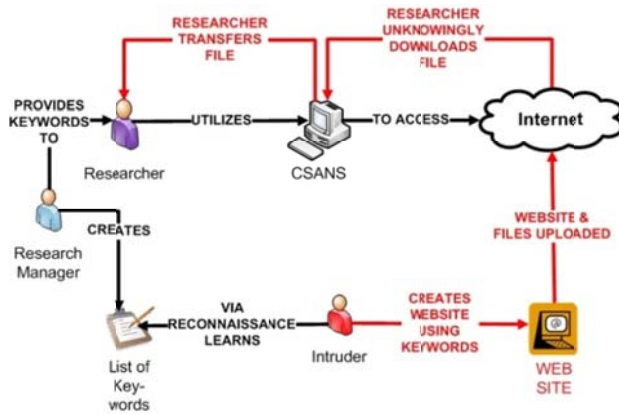


Figure 3: Interactions between key players and entities in experiment

Specifically, Figure 3 shows that the research manager creates a list of keywords and provides them to the researcher. From there the researcher utilizes the CSANS to access the Internet. Meanwhile an intruder has learned about the keywords and content associated with the target organization. With this knowledge the intruder creates a website that is loaded with both non-malicious and malicious files that are capable of both tracking and infecting the researcher’s systems.

The researcher via a simple search locates the website and believes it to be a “treasure” of information related to the task. As a result, the researcher could potentially download files that may yield information of interest to the intruder.

1.1 Ethics and Institutional Review Board for Human Subjects

Given the sensitivity and potential ethical issues associated with this type of experiment some organizations and its users may feel that purposely trying to compromise systems and discover PII is a violation. However, we were authorized to conduct this experiment against the CSANS users with the support of the internal review board and the permission of upper management.

To address any ethical issues that could stem from a study that involves humans, we requested the permission of the organization’s review board, who subsequently authorized this experiment to ensure that the CSANS and its users were following best practices when it came to protecting the anonymity of the organization.

One of the major difficulties we encountered was to prevent people not associated with the experiment from accessing our website and downloading malicious files on to their system. We considered putting up a firewall to block access to our website except from a certain range of IP addresses. This proved to be difficult as the CSANS IP address changes randomly on a daily basis as part of the anonymization process.

Next, we considered password protecting our website but were concerned about negatively impacting the experiment. This is because it would not have made sense for the researchers to know the password associated with our experimental website. However, we felt that this was the best option and worth preventing the risk of adversaly infecting unsuspecting users.

Specifically, we password protected the malicious files from Table 3 and gave the password to the research manager. The research manager then created a fictitious story as to how he learned the password and passed it on to his researchers as part of the task. While not a perfect solution, we did attempt to prevent people who accidentally visited our site from infecting their machines. We leave it to future experiments and research to come up with a more eloquent solution to this problem.

1.2 Roles and Responsibilities

In Table 1, we describe the roles and responsibilities of several key people who participated in the experiment.

Table 1: Roles and responsibilities

Roles	Responsibilities
Research Manager	<ul style="list-style-type: none">• Authorizes the experiment;• Creates and provides the list of keywords to the researchers;• Manages the researchers;• Authorized to stop the experiment;
Information Security Manager	<ul style="list-style-type: none">• Develops, secures, and monitor the CSANS;• Assists with security related issues;
Author of Paper	<ul style="list-style-type: none">• Acts as the intruder;• Develops the experiment;• Conducts the experiment;• Creates the website and corresponding files;• Captures and analyzes the data;• Documents and reports on the findings of the experiment;
CSANS Users (Researchers)	<ul style="list-style-type: none">• Does not have knowledge of the experiment;• Completes research task independently with no collaboration;
CSANS Organization	<ul style="list-style-type: none">• Has no knowledge of the research task or experiment;

2. Background Information

In this section, we briefly discuss anonymity on the Internet. Depending on the individual reason for accessing the Internet, anonymity may or may not be an important issue.

For example, anonymity requirements may not be a concern for someone who is simply using the Internet to access a news site. Similarly, there is a growing trend for people willing to disclose their identity by posting personal information about themselves on a social networking site.

Conversely, users in the defense and intelligence related fields almost always have a requirement to remain anonymous when researching and compiling information from other (often hostile) countries; if their identity is known it makes completing the research far more difficult and possibly even dangerous. For example, the groups and countries that they are researching could potentially retaliate in response to the research being conducted without their consent. For the purposes of this paper, the CSANS organization consisted of real users who needed to remain anonymous in order to complete the research tasks that they were given.

2.1 Anonymizing Network Systems

One popular way to help protect a user's identity against a possible intruder is to implement an ANS. The field of anonymous communications and specifically, ANSs started in 1981 with David Chaum's Mix [3]. An ANS is a system that is designed to protect a user's identity while they are using a computer system that can access the Internet. In general, an ANS is a tool that attempts to make activity on the Internet untraceable.

The first goal of an ANS is aimed at preventing the true identities of specific hosts from being leaked such that an audit trail of user activity cannot be formed. The second goal is to prevent the true identities of internal hosts from being leaked such that a map of supported services can be constructed. The third goal is to prevent the leakage of specific security practices within the publishing organizations network.

There exists a plethora of different ANSs including the aforementioned Anonymizer and Tor that are available on the market today. Other ANS solutions exist and they include but are not limited to: Browzar [20], JAP [21], and SafeSurf [22].

2.2 Inner-workings of the CSANS

The CSANS is a variation of Anonymizer from Anonymizer, Inc [1]. The CSANS has been designed to prevent an open Internet connection from being exploited. All traffic is routed through dedicated hardware housed in a secure facility. Only authorized administrators are given access to this facility. The CSANS maintains a large pool of IP addresses and is able to easily rotate and replace IP addresses that have possibly been compromised. The

users of CSANS should expect to see their IP address change about once every 24 hours. This concept is known as the IP rotator scheme.

In this scheme, anonymized users are mixed with “regular” worldwide consumers. The behaviors of all users (both anonymized and regular) are mixed together in an effort to try and prevent behavioral patterns from being noticed. The regular users are diverting attention away from the anonymized users by virtue of the search queries and traffic that they are generating.

The hardware associated with CSANS is essentially an enhanced version of the Juniper Virtual Private Network (VPN) box. The CSANS includes two interfaces, the front-end and the back-end. The front-end interface has a specific network IP address and the back-end has a completely different network IP address. For security purposes, the IP addresses are visible internally but not externally. The firewall routing for CSANS is designed to protect users from accidentally visiting sites associated with the specific organization that implemented CSANS. The routing rules are set to essentially mask the source and destination as to provide a second layer of anonymity.

The CSANS is hosted on a Windows server virtual machine and users of the CSANS will have to remote desktop into the system to get access to CSANS as seen in Figure 4.

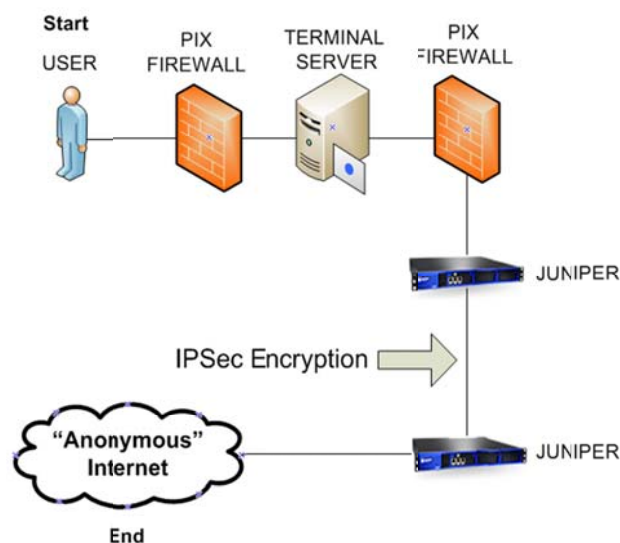


Figure 4: CSANS architecture

The CSANS environment is locked down so that only certain users have administrative privileges.

Since this particular implementation is essentially a shared environment, all users’ requests, including software installations will need to be authorized. This follows best practices for the security aspects of the server. Specifically, in an effort to keep the CSANS clean from malicious code, the system administrators incorporate the use of a virtual machine snapshot. Once a week anti-virus updates, Windows updates, and new software installations are performed resulting in a new snapshot being taken. If the current snapshot becomes corrupted or infected with malware, CSANS would be able to revert back to a known good snapshot.

One of the key features of CSANS is the ability to transfer files from CSANS back to the client OS. While this is extremely helpful to the end-users in terms of usability and functionality it is also a possible security concern. The rationale for implementing the transfer file component is because the CSANS users have a requirement to share files and to create and analyze reports using software from their client machines.

The developers of the CSANS implemented a simple file sharing system. The file sharing system utilized an FTP client on their machine along with an FTP server that could run from the CSANS environment. The CSANS users simply saved their files to a “Home” directory and using FTP moved the files to their client machine.

2.3 Targeting and Profiling

In this section, we discuss targeting a specific organization and creating a document to be used in the attack. The goal of the intruder is to get the target to download the malicious document. To do this, the attacker needs to profile the target so it can create documents that are appealing. Ideally the documents that are uploaded on the website will be related to a topic that is of interest to the user. In addition, most users feel comfortable with Adobe PDF and Microsoft Word documents feeling that it is safe to open these types of files.

Therefore, if the intruder can customize the document to make it look like it was an organizational newsletter, job posting, company form, etc; the more likely it is that the document will be downloaded [5]. If intruders can build up a good data set it would only be a matter of time before an employee comes to the intruder’s website downloads something malicious.

The attack that we will demonstrate is based on the idea that the intruder can use “common and public” knowledge about a particular organization and its users. If, for instance, the intruder can infer and attract attention about the target it will make it more likely that the attack will succeed. Through profiling we were able to identify the CSANS organization as one that is related to military, defense, and intelligence. Therefore, we create documents that were related to these same areas of interest.

2.4 Survey of Current Research on ANS Attacks

Most users of ANSs gain a false sense of security as they are under the assumption that they are completely protected and anonymous. Almost all ANSs including the one used for this case study are vulnerable to a variety of different types of attacks.

In this section, we describe current research that is being conducted to attack ANSs. The literature review considers traffic analysis, application, and network-based ANS attacks.

To begin, we consider a work titled “*Performance Analysis of Real Traffic Carried with Encrypted Cover Flows*” [6]. The authors discuss the idea that network encryption, both at the packet and session layer, is used widely for securing private data. They use simulation and an analytical model to examine the impact on user experience via a scheme that masks the behavior of real traffic by embedding it in synthetic and encrypted cover traffic.

In the article “*Playing Devil’s Advocate: Inferring sensitive information from anonymized network traces*” [4], the authors attempt to solve the problem of publishing data that can potentially leak sensitive information about the publishing organization. The article introduces the problem by suggesting that it is imperative that trace and log data be made publicly available for verification and comparison of results. The authors attempt to conduct an analysis and create techniques to infer sensitive information from these network traces. The study as performed by the authors demonstrates that there are more substantial forms of information leakage that inherently compromise current anonymization methodologies.

In “*Network Flow Watermarking Attack on Low-Latency Anonymous Communication Systems*” [9], the authors describe the concept of watermarking the

network traffic in an effort to break the anonymization supplied by the ANS. The watermarking aspect in this case can potentially allow an intruder to influence the traffic thus allowing them to discern the identity of the user. The authors were able to penetrate the Total Net Shield, the “ultimate solution in online identity protection” of www.anonymizer.com, which is almost exactly the same ANS that is used in the case study associated with this paper [9].

In “*On Web Browsing Privacy in Anonymized Netflow*” [10] the authors improve on previous research done on reconstructing web browsing activities from anonymized packet-level traces. This is accomplished by accounting for real-world challenges such as browser caching and session parsing. This research evaluated the effectiveness of the author’s techniques by identifying the front pages of the fifty most popular websites on the Internet.

In the article, “*How Much Anonymity Does Network Latency Leak?*” [11] the authors present two attacks on low-latency ANS schemes. The first attack allows a pair of colluding websites to predict, based on local timing information and with no additional resources, whether two connections from the same Tor exit node are using the same circuit with high confidence. The second attack requires more resources but allows a malicious website to gain several bits of information about a client each time a user visits the site. The authors evaluate both of their attacks against the Tor network and the MultiProxy proxy aggregator service.

In a related article, “*Peering through the Shroud: The Effect of Edge Opacity on IP-Based Client Identification*” [12], the authors state that online services often use IP addresses as client identifiers when enforcing-access control decisions. The idea of the paper is to determine the impact of edge technologies such as NAT, proxies, and DHCP have on the utility of using IP addresses as client identifiers. Specifically, whether the IP address of an incoming client is a useful identifier of access control decisions and to explore the extent to which edge opacity obscures a server’s view of the client.

In the article, “*Taming the Devil: Techniques for Evaluating Anonymized Network Data*” [13] the authors’ primary concern is to evaluate the efficacy of network data anonymization techniques with respect to the privacy that they afford [13]. Specifically, the authors are trying to make the

network flow uniquely identifiable even after it has gone through the ANS. They are also considering techniques for evaluating the anonymity of network data and to simulate behavior of an intruder whose goal is to deanonymize the objects (host or web pages) [13].

In the article, “*Privacy Oracle: A System for Finding Application Leaks with a Black Box Differential Testing*” [14], the authors described the design and implementation of a system called Privacy Oracle that is capable of finding application leaks. The authors attempt to solve the problem of application leaks by using black box differential testing. The creators of Privacy Oracle are most interested in what information was leaked, when it is exposed, and who can receive it.

In an article titled “*Abstracting Application-Level Web Security*” [15], the authors investigate new tools and techniques which address the problem of application web security. The authors describe their solution to address the problem of application layer web security by describing a scalable structuring mechanism facilitating the abstraction of security policies from large web applications developed in heterogeneous multi-platform environments [15]. The second aspect to their solution is to represent a tool which assists programmers developing secure applications which are resilient to a wide range of common attacks. Finally the third aspect is to report results and experiences arising from the implementation of these techniques [15].

Based on a review of current literature we found that many security researchers are studying attacks that focus on traffic analysis, network-layer based attacks, and some application-layer based attacks. In this section, we introduced a few different anonymization schemes and concepts including ANSs, network trace anonymization, and to a lesser extent proxy systems. While each of these systems have unique properties and disjoint security goals we felt compelled to provide an overview of the current state of attacks against ANSs. We introduced de-anonymization attacks that are complex and detailed in nature when compared to the attack that we are demonstrating. Again what we are suggesting is something far simpler and that is users and certain applications, if successfully compromised, can expose the identity of a user behind an ANS.

In fact, ANSs are not designed to prevent against basic attacks against applications and the users. The attacks we describe are quite effective in terms of

how they can be used to produce client information such as IP addresses, email addresses, source location, timestamps, and other pertinent system information.

2.5 Google Location

While not an attack by definition, Google Location feature has a major impact on operation security associated with ANSs users. Specifically, the Google Location technology in certain cases identifies the ANSs users of their true location [18]. Google has massive amounts of technology and tools to be able to analyze traffic and find correlations between the user’s search criteria and the user’s likely locations.

Essentially, users forget or are unaware that they are performing a search that should not be done while on an ANS such as CSANS. This concept often referred to as tradecraft is when a user searches for local businesses or perhaps performs search queries including but not limited to local politics, weather, and sports. As a result, Google is often able to correlate these searches and show the user’s real location regardless of whether they are behind an ANS.

For example, suppose that the CSANS IP address is 123.456.789 and it shows that the location is anonymized to be in Ghana. Now let us assume that a few users make the mistake of performing attributable searches using Google on CSANS such as searching for local restaurants in their real home city of Atlanta, GA. Google is then able to identify this behavior and given enough searches and traffic will assume with a high likelihood that this IP address is not associated with Ghana but instead with that of Atlanta, GA. The user’s behavior is being used by Google to find the “most likely” location from which this user and corresponding IP address is coming. Given enough of this data the Google Locator actually becomes quite accurate.

This problem is not necessarily specific to Google although they are by far the biggest actor in developing this technology. It should also be pointed out that Google is not “attacking” in an effort to be malicious or uncover the true identity of any organization. Instead they are simply trying to be more precise with their search results and advertisements. From Google’s perspective they want to be both complete and accurate with their search results.

2.6 Intruder Threat Model

Threat modeling is a method of assessing and documenting the security risks associated with an application. The threat model for CSANS is based on the process outlined in [7].

First, we examine what PII the intruder is after. The intruder is most interested in the true identity of a user behind the CSANS. Such information could include:

- User’s real name
- User’s place of employment (organization name)
- Real IP address (non-anonymized IP address)
- Browser configuration and browser history
- Sensitive documents that reside on users system
- Software and programs currently installed on CSANS
- Configuration and implementation of CSANS

One method for identifying and categorizing threats is known as STRIDE. This method is a classification of the effects of realizing a threat and it stands for Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege [7]. The STRIDE classification for various vulnerabilities is listed in Table 2. Next we consider internal and external threats, the following are a few threats associated with the CSANS.

- The CSANS users have been known to post sensitive documentation to fake email accounts for the purposes of storing data for future use.
- The CSANS administrator has full control over all the user accounts and passwords.
- The CSANS system itself has directories of sensitive documents that are stored on the server itself (before it is transferred). By logging into the server an intruder could view these documents.
- The CSANS has a built-in transfer system that is designed for convenience rather than security.
- The anonymization is being done by a third-party.

- Users have been known to perform attributable search queries while on the CSANS.

In addition, we included a DREAD rating for each of the CSANS vulnerabilities that we identify in Table 2. The DREAD rating is a method for characterizing the risk associated with vulnerabilities. The DREAD rating comprises of Damage potential, Reproducibility, Exploitability, Affected users, and Discoverability [7].

The DREAD rating for the vulnerabilities associated with CSANS, used a three-point rating scale. A rating of one indicates that the impact is minor, the number of systems impacted is a small percentage, and the likelihood that the attack is identified is low.

In general, the lower score equates to minor damage, less reproducibility, less exploitability, fewer numbers of affected users, and less likely to be discovered. Conversely, a rating of three equates to major damage, more reproducibility, more exploitability, larger numbers of affected users, and more likely for the attack to be discovered. The final DREAD rating score is simply the average of the scores for each of the five aforementioned categories.

Table 2 displays a non-exhaustive list of several vulnerabilities that the CSANS intruder could take advantage of along with the corresponding DREAD ratings.

Table 2: Vulnerabilities associated with the CSANS

Vulnerability #1	Intruder is able to determine browsing patterns. [See Section 3.5 for more details of this vulnerability]
Description	An intruder inspects the traffic as it goes through the server and onto the anonymized network.
STRIDE Classification	Tampering Information Disclosure
DREAD Rating	Damage Potential = 3 Reproducibility = 1 Exploitability = 3 Affected Users = 3 Discoverability = 2 DREAD Rating = 2.4

Vulnerability #2	Intruder attacks the de-anonymization process.
Description	The intruder attacks the de-anonymization servers and/or the processes associated with the endpoint.
STRIDE Classification	Tampering Information Disclosure
DREAD Rating	Damage Potential = 3 Reproducibility = 1 Exploitability = 2 Affected Users = 1 Discoverability = 1 DREAD Rating = 1.6
Vulnerability #3	Intruder performs any number of the attacks outlined in section 2.4.
Description	By performing any number of the attacks outlined in section 2.4 the intruder may be able to de-anonymize CSANS.
STRIDE Classification	Tampering Information Disclosure Denial of Service
DREAD Rating	Damage Potential = 2 Reproducibility = 2 Exploitability = 2 Affected Users = 2 Discoverability = 2 DREAD Rating = 2.0

3. Methodology

The following sections describe the experimental methodology. Specifically, we discuss the experimental goals and objectives, experimental setup, and the creation of a fictitious research task and website.

3.1 Experimental Goals and Objectives

The primary goal of the experiment was to determine the extent to which we could use application and user based attacks to discern the true identities of the CSANS users. Also, we wanted to show that network level anonymization is not enough and that there needs to be better education to promote better anonymization. To accomplish this goal we created documents that were used to both track and infect the CSANS users. We also

examined the impact of users performing attributable search queries while on CSANS. We relied on analysis of traffic that we captured to find key points of interest.

3.2 Experimental Setup

The experiment for our case study lasted for two months at the request of the CSANS research manager. There were twenty CSANS users that were included in the experiment and they were asked to conduct all research within the confines of the CSANS.

Prior to the start of the experiment, we created and hosted an inconspicuous website and ensured that both the non-malicious and malicious files were tested before being uploaded. Furthermore, we ensured that file logging and tracking was enabled to allow us to capture statistics from the web server. The website utilized various sources for content and was approved by the research manager as being enticing enough to attract the researchers.

The next step was to ensure that the website could actually be found via a simple search. By using Google search engine optimization and similar search ranking tools we found that our website could in fact be located by virtue of simple search logic. A major aspect of the proposed study design was to determine the attack vector. For the purposes of this experiment we focused exclusively on Adobe PDF.

The design of our study focuses on three major aspects, namely impact, stealth, and depth. The impact aspect measures how many CSANS users actually accessed the website and clicked on the files. The stealth aspect indicates how many users observed and notified management about the website and/or malicious file. Lastly, the depth aspect quantifies the quality of the malicious file as it relates to the amount and value of PII that a potential intruder could discern.

3.3 Fictitious Research Task

The CSANS users were specifically asked to complete this research task that focuses on Iran, military, defense, intelligence, and counter-terrorism issues. This fictitious research task is like any other task that the CSANS users would typically investigate.

To begin the experiment, the research manager provided the CSANS users with a list of keywords to search on as part of the task. The keywords include:

Quds Force, Bank Markazi, Hezbollah, Zahedan, Herat, Opium, Taliban, Battery Acid, and IRGC.

The assumption here is that the keywords are general enough to be known by a potential intruder as part of their targeting and profiling. Therefore, simply protecting the keywords is not a viable security measure.

3.4 Website

The necessary pre-requisite for this experiment was to design and build the aforementioned website. We researched each of the keywords independently and extracted content from other related websites to build into our own custom website.

Since the goal of the experiment is not a demonstration in website design, we developed a basic HTML website and hosted it on a third-party hosting site. The website corresponded to the different topics associated with the keyword list and using simple logging techniques, the website can log the actions of the users.

The files were uploaded to the “documents” section of the website. The files included documents associated with the keywords. Each file was randomly selected to fall into one of three categories. The first category was that the files were original and completely clean. The second category was that the files were embedded with a Javascript (app.launch) function that would automatically open up a new webpage in the user’s browser. For example, the file would open a webpage called www.website.com/3490 where 3490 is a random number that can be easily identified in the data logs. The third category was that the file itself was embedded with malicious code via the Metasploit Framework (see Section 3.5). If the malicious code is opened in a vulnerable version of Adobe Reader on a vulnerable Windows OS, the user’s machine could possibly become compromised. Once compromised it would be trivial for the intruder to de-anonymize the user. Table 3 shows a sample summary of the document name, the corresponding category, and a description of the exploit where applicable.

Table 3: Sample summary of document and action

Document (.pdf)	Category	Action
Iran Banking	#1 Clean	Downloads and opens
Herat	#2 Redirection	JavaScript app.launch to random page on website
Jihad Encyclopedia	#2 Redirection	JavaScript app.launch to random page on website
Taliban Poppy	#3 Malicious	Redirects to webpage being listened to by Metasploit: multi/handler listener
Iranian Affairs	#3 Malicious	Infected w/ “Aurora” exploit
Iranian Jobs	#3 Malicious	Infected w/ “Adobe (Collect_Info)”

3.5 Metasploit

In order to create the websites malicious files we relied on the Metasploit framework [8]. The Metasploit framework was built to provide useful information and tools for penetration testers, security users, and intrusion detection system (IDS) signature developers. The Metasploit framework was created to provide information on exploit techniques and to create a functional knowledge base for exploit developers and security professionals.

The intruder would utilize the Metasploit framework to create attacks that target items such as specific applications, operating systems, and web browsers. The exact steps of finding the exploit, setting/executing the payload, and listening are outside the scope of this paper.

3.5 Capturing Traffic

We also consider a second threat model associated with the intruder’s ability to capture traffic which is shown in Table 4.

Table 4: Threat Model: Capturing Traffic

Source of Traffic (Example)	Success	Attack Level Success Rate
	Complexity	
ISP (Comcast, Cox)	Low	Low
	High	
Service (Facebook, Gmail, Amazon)	Low	Medium
	High	
Custom Website	Low	High
	High	

In the ISP case, the anonymized traffic blends in with the traffic from regular traffic from subscribers to that particular ISP. Assume that the ISP is Comcast or Cox Communication. There are tens of thousands of users that are accessing the Internet by virtue of these ISPs. In addition, there are several search queries on any number of topics that are being performed and this traffic will blend together.

As it stands, the attack level is very low because it would be difficult for an adversary to determine if the search queries and the traffic were from a truly anonymized user or a “regular” user who coincidentally was searching for the same information. Also, it is more difficult to obtain access to a big ISP and monitor all communications. The complexity and cost for doing so is high when compared to poisoning Google or posting an advertisement to CNN.

For example, consider the CSANS user who is researching Iran and comes across a website on the terrorist group “Hezbollah” and finds it pertinent to the research task. Now suppose a student who has Cox as their ISP comes across the same website for a high school research paper that they are doing. From an attackers point of view it would be very difficult to determine whether either or both the CSANS users and the high school student were attempting to hide their identity given only the traffic obtained from the ISP.

In the service (external website) case, the anonymized traffic blends less when compared to that of the ISP traffic. When someone goes to an external website such as Facebook, Gmail, or Amazon they give away some of their identity by virtue of logging into a site, accessing personal

items, and cookies (e.g. Amazon book preferences). If this traffic is obtained by an attacker the success rate of determining the identity is increased when compared to simply obtaining the ISP traffic.

For example, consider a CSANS user who accesses their personal Facebook account. The CSANS user must login to Facebook, perhaps views photos, sends a message via Facebook, and posts a few messages. There is something inherently personal about the example that I just gave. An attacker who is able to acquire the traffic can often determine the true identity of the user as we will show in the results section of this paper.

In the final case, as was described earlier we created a custom website that is meant to entice the CSANS users. The website content is associated with a list of keywords that a typical CSANS researcher would use as their search query. Included in this website are a series of file that serve different purposes. The files have the ability to track the users and to also infect the users. The files if downloaded (and transferred) to their client machine will potentially give the adversary the ability to discern the user’s true identity. It is for this reason that the success rate of the attacker is the highest. We will show the success rate of this novel idea in the results section. The complexity of obtaining traffic from a custom website is also low because there is no dependence on IPS or popular services.

To substantiate our claims, we also attempted to gather and analyze the traffic that was generated by the researchers while on CSANS. We ran the Wireshark packet analyzer to accomplish this [23]. The generated Wireshark .pcap was then run against a tool called tcpextract and used to carve headers and footers from the traffic [24]. This extraction yielded several web images, cookies, files, and portions of web activity.

Next, by using a custom developed Perl script we were able to take the destination IP address that was captured and in most cases identify the hostname, internet registries (e.g. apnic, arin, ripe), autonomous system (AS) number, the country where the servers are located, the routing information, and the date the website was registered.

We then took all of this information and created a very simple Microsoft Access database so that we could perform queries on the data that we found. The database included over 10,000 records as a result of running the experiment and capturing the traffic. The following section describes our findings.

4. Results

Based on an analysis of the traffic that was generated by the CSANS researchers we were able to identify a variety of data points that could have been used to identify the CSANS. First, we saw several connections being made from the real (un anonymized) IP address of the CSANS. The source of the IP address was most likely as a result of the researchers utilizing remote desktop (RDP) from their client (un anonymized) system to access CSANS.

Furthermore, we identified almost daily requests to Facebook. This was a concern because of the obvious attributable nature of this social networking website. We then began to analyze traffic, cookies, and ultimately uncovered two distinct email addresses via the Firefox profiles. Note: It was necessary for us to download SQLite Database Browser in order to extract this information from Firefox [25]. Going one step further, we logged on to Facebook and did a “*Friend Finder*” search based on the two emails that we uncovered. Sure enough, the results of our search yielded the names of two different people that were later shown to have been part of the research group at CSANS.

Next we discuss the concept of Facebook profile ID. For example, when you login to Facebook (assuming you have not replaced id with username) you will see something similar to the following whereby the ## correlate to real numbers:

<http://www.facebook.com/profile.php?id=604###>

During our capture we were able to identify two distinct Facebook profile ID’s. However, when we accessed the profiles we were surprised to see that it did not relate to any CSANS users. In fact, it was the profiles of two different men living in Iran who were not associated with the CSANS organization whatsoever. We still need to perform future research to understand our findings with respect to profile ID. However, it seems to indicate that the CSANS researchers were searching for these two individuals as part of their research task. This has not yet been confirmed or denied. For more information on Facebook ID the reader is urged to view [26].

The traffic analysis also uncovered the researchers accessing amazon.com, continental airlines, and several other US based companies. Furthermore, there were other attributable websites

such as LinkedIn and Twitter that were being regularly accessed by the CSANS researchers.

By using the tcpextract tool, we were able to uncover portions of email that were created from Gmail and Thunderbird email client. While it was difficult to view the email in its entirety it appeared to have a subject associated with US military.

By capturing the traffic for such a prolonged amount of time we were able to find various patterns associated with the time of day users log in to CSANS, the types of sites (both attributable and non-attributable), and the routing of the traffic as it makes its way to these sites. We then compared this traffic to that of traffic acquired from a University in Greece (Note: One of the co-authors is from Greece and was able to acquire this traffic for the sole purposes of this experiment) and found that it was quite easy to tell which traffic patterns were associated with what entity.

In addition, to capturing the traffic and utilizing tcpextract, we also made use of a tool called network miner [27]. This tool allowed us to import a Wireshark .pcap file and mine the traffic. The tool quickly places the traffic into categories such as hosts, frames, files, images, messages, credentials, sessions, DNS, parameters, keywords, clear text, and anomalies. The main results that were acquired from this tool were the search criteria and queries used by the researchers. For example, we found the following search criteria used by the researchers “*Iran petrochemical production process*” which is a perfectly acceptable search query given the research task. However, we also came across a search query for “*Fairfax County School Closings*” that is clearly more attributable to the CSANS location.

In the actual fictitious website experiment, we saw the anonymized IP address 207.195.x.x that was identified as Global TAC, LLC. This log entry showed up several times throughout the experiment on any given day. We identified that several of the files from the website were downloaded by this IP address. Shortly thereafter, we could see that our random page was identified in the log files but this time with the NAT IP address which was confirmed to be that of the target organization. This was repeatable as we saw new IP addresses being related to the CSANS IP address. Also, most of this activity appeared during 8:00 AM EST and 5:00 pm EST, the timeframe most US-based organizations are open for business.

In our logs we were able to clearly see the web browser and the operating system of the user who opened a certain page in our website. This information could be extended to other de-anonymization projects such as [17]. The intruder could also use this information to customize the Metasploit exploits.

In addition to the logs supplied by our web server we also utilized tracking statistics from [19]. This allowed us to have a separate tracking mechanism in place and we were able to identify on several occasions that a user was in fact logged into our website. We also noticed that both the CSANS anonymized IP address and the case study agency were logged in at the same time. Since the traffic to our website is light it became quite clear that there was some relationship between the two sets of IP addresses that are simultaneously logged in. This was further supported by the fact that both users logged out of our website at nearly identical times.

Within the logs we were able to determine whether or not we have seen this “user” before and in some cases determine the first time they visited and the last time they visited. In addition, we were able to identify the system hardware used and the browser, browser language, operating system, and screen resolution. Next, we could uncover the ISP and NAT IP address of the CSANS organization. Another piece of data that we saw in the logs was the hostname of the CSANS. This hostname appeared to be fictitious as it was the name of a Star Trek character. This information was then looked up via *nslookup* and came back with an IP address that is tied to the CSANS web proxy. Finally, we were able to identify the connecting city, state, country, timezone, and latitude/longitude.

One particular user was paying close enough attention and identified a suspicious “tracking code icon” that was intentionally displayed on the site. However, as far as we can tell this was the only time that anyone from the CSANS identified the website or the experiment as suspicious.

Also, we were able to see that on certain days (due to what the IP address was on that given day) that Google Location was able to identify the location of the CSANS organization within about 20 miles. However, on other days the Google Locator was not able to identify the CSANS location. These results show that at least one CSANS IP address was compromised. This was likely due to attributable search queries being performed. It is not immediate

clear whether or not it was a CSANS user that was responsible for the attributable search queries.

Lastly, one user downloaded a malicious file, transferred it to their vulnerable Windows XP computer, and opened it in a vulnerable version of Adobe Reader. By utilizing the Metasploit multi/handler listener we were able to insert the Metasploit payload and get a shell on to the CSANS users system. At this point, we could easily have dropped a keylogger, taken a screen capture of the page, run system commands, and quite easily harvested information that would result in determining the user’s true identity.

5. Conclusions

The results of our traffic analysis and experiment were conclusive in determining a few unique CSANS users. In general, the results of the experiment showed that the anonymized IP address could be directly correlated to the CSANS organization. This was shown by looking at the logs and cross referencing a series of random IP addresses with the CSANS organization’s NAT address. This information was most likely due to the user downloading the file, transferring it to their computer, opening the file, allowing JavaScript to run thus causing a random page within the intruder’s webpage to launch. In the log files, it was easy to see that a given page of our website was accessed.

Specifically, we were able to conclude that users who transferred the malicious document to their client operating system were at greater risk. This is because the client operating systems and applications were not nearly as secure as the CSANS. In some cases, the client operating system was unpatched and utilizing an obsolete version of Adobe Reader.

Furthermore, we were able to conclude to some degree that attributable searching on CSANS did play a role in determining the location as shown by Google Location. It is not known at this time what the specific search criteria used was, nor is it known whether or not it was a CSANS user that performed the attributable search. We were able to show that given a clean CSANS virtual machine snapshot, certain IP addresses used in the CSANS IP rotation scheme were identifiable as coming from a location in close proximity.

Overall, we were able to show that we can break the CSANS anonymization without doing anything

sophisticated. Instead we relied on targeted attacks that focused on the weaknesses of the users and applications.

Based on the results and the conclusions, we recommend several possible defenses that the organization, the CSANS, and the users could have implemented to prevent the attacks.

One possible defense would be to reconsider the use of the transfer file system. Even though the transfer file system is a convenient way to move files it is one of the weakest links in the CSANS security architecture. At a minimum, it is recommended that the users create new and clean PDF files from the ones that they download. If the CSANS organization decided to keep the transfer drive it is highly recommended that the files be run through multiple anti-virus engines and/or a sandboxed system that could execute the file, trigger on illegal system operations, and therefore potentially block the exploited file from connecting back to the intruder's listener server.

The transfer drive system can be extended with a secure and dedicated analysis platform. An improvement to the system architecture might allow the transfer of files to another virtual machine that would allow for execution but prevent information flows to untrusted hosts. This scheme, similar to what has been done with honeypots, would address the vulnerabilities in the current CSANS, while at the same time, allowing for richer analyses to be performed.

One simpler defense is to patch and harden the client machines. However, we believe that the intruders will continue to create exploits and eventually will be able to infect the client machines. This defense will reduce the risk but not completely prevent it.

The users should be cautious with allowing Javascript to run on their computers. Options such as noscript from Firefox [16] will automatically prevent JavaScript from running. The users should have immediately contacted the organization's information security group when their browser started to open up new browser windows.

To defend against attributable searching while on the CSANS the users need to be cognizant that logging into personal accounts such as email, social networking, and banking sites could provide information that could discern their identity. Next, the users are advised not to perform any searches associated with their true geographic location. For

example, if the user is located in Denver, CO they should refrain from searching for items such as Denver sports teams, weather, traffic, jobs, and restaurants. This information can be put together to yield the users identity.

We demonstrated in section 2.5, Google is already proving this by virtue of their Google service location feature that is tied to a specific IP address regardless of whether that IP address is anonymized [18]. For this reason, we have recommended that the CSANS users complete all of their searches using www.google.com.gh whereby the "gh" is the country code for Ghana. We chose Ghana primarily because it is an English speaking country that is often attributed to the IP anonymization scheme that is used at the CSANS. Still another possible defense would be to set the Google Location to be a random place in the country. This concept is commonly referred to as "artificial pinning" and is just another small measure that could potentially help to prevent the identity of that particular IP address from being known. At the time of this writing, it was not possible to change the location from outside of the United States assuming that you were using [www.google.com\(.en\)](http://www.google.com(.en)) as your main search engine. Also, in some cases it is not even possible to change this setting and/or it may not appear for a plethora of reasons that are outside the scope of this paper.

Similarly, the users should not provide any personal information to a website, toolbar, or widget about their true location when logged on to CSANS. For example, users often set weather monitoring tools to be their local zip code. The CSANS developers should create visual guidance when a user is using the non-attributable CSANS to search for information related to their true location. Perhaps a flashing notification or a confirmation screen before a given search is executed would be useful.

Another solution is to provide more non-anonymized users to share the IP address space of the anonymized users. This will create a large quantity of search results that could allow the anonymized users to potentially hide behind an overwhelming amount of search data. For example, if instead of using the CSANS the organization tunneled a portion of their traffic to universities and colleges in the same area it might actually improve the anonymization. This is because a potential intruder would find it difficult to determine the source of the traffic and whether it was anonymous.

Since the research being performed by the CSANS organization might be similar to that of a university professor it would be nearly impossible to determine which one is anonymous.

We conclude this paper by looking at future research. This experiment was done on a very small scope with little money available to the authors. Future research could extend this experiment and consider different anonymization schemes such as Tor and different applications such as Microsoft Office and Adobe Flash. We leave it to future research to determine if attacks against other ANS users would be successful. Also, this experiment was narrow with a focus on military, defense, and intelligence. It would be interesting to determine if a similar attack would work in other organizations such as healthcare or banking. Future research could provide methods to improve the defenses offered in this paper and more importantly help educate the users on how to best defend themselves from the three main attack methodologies.

While not a new idea, users remain the weakest link in any non-trivial security scheme. The implication, however, is the researchers were careless or ignorant of basic security principles. The fact that they betrayed personal information so easily does not reflect well on the level of security education provided to the researchers.

One of the main takeaways from this study was that both the administrators and users of the CSANS were not sufficiently educated about quite prevalent attack vectors for compromising client systems and violating user privacy.

To address this problem, we briefed the CSANS researchers on the results of our findings. We also explained to them the dangers of accessing personally identifiable websites while on the CSANS. The training that we performed was over a two day span and it we discussed our findings along with our recommendations as to how to resolve the problem. We are planning on conducting a similar experiment in the upcoming months in order to measure the effectiveness of our training with regards to combating the attacks mentioned in this paper.

This paper illustrates that regardless of the anonymization scheme a given organization has in place the users are often the weakest link. It is recommended that users be trained on common pitfalls when it comes to downloading documents and searching the Internet while on an ANS.

References

- [1] Anonymizer. URL: <http://www.anonymizer.com>
- [2] Tor. URL: <http://www.tor.net>
- [3] D. Chaum, "Untraceable electronic mail, return addresses and digital pseudonyms", *Communications of the ACM* 24, ACM, New York, NY, February 1981, pp. 84- 90.
- [4] S.E Coull, C. V. Wright, F. Monrose, M. P. Collins, and M. K. Reiter, "Playing devil's advocate: Inferring sensitive information from anonymized network traces," in *NDSS: Proceedings of the Network and Distributed System Security Symposium*, 2007.
- [5] HD Moore, V. Smith, "Tactical Exploitation", Black Hat USA 2010, Course Slides, Las Vegas, NV, July 2010.
- [6] N. Schear and D. M. Nicol, "Performance analysis of real traffic carried with encrypted cover flows" in *22nd Workshop on Principles of Advanced and Distributed Simulation*, 2008.
- [7] F. Swiderski, and W. Snyder, *Threat Modeling*. Redmond, WA: Microsoft Press, 2004.
- [8] Metasploit. URL: <http://www.metasploit.com>
- [9] X. Wang, S. Chen, and S. Jajodia, "Network Flow Watermarking Attack on Low- Latency Anonymous Communication Systems," In *Proceedings of the IEEE Symposium on Security & Privacy (S&P)*, May, 2007.
- [10] S.E Coull, M.Collins, C. Wright, F. Monrose, and M. Reiter, "On Web Browsing Privacy In Anonymized Netflow," *Proc. Of the 16th USENIX Security Symposium*, Aug., 2007.
- [11] N. Hopper, E. Y. Vasserman, and E. Chan-Tin, "How much anonymity does network latency leak?," *ACM Transactions on Information and System Security*, forthcoming 2009.
- [12] M. Casado and M.Freedman, "Peering through the Shroud: The Effect of Edge Opacity on IP-Based Client Identification," In *Proceedings of the 4th*

Networked Systems Design and Implementation, April, 2007.

[13] S.E. Coull, C. Wright, A. Keromytis, F.Monrose, and M.Reiter, "Taming the Devil: Techniques for Evaluating Anonymized Network Data," *In Proceedings of the Network & Distributed System Security Symposium*, San Diego, CA, Feb., 2008.

[14] J. Jung, A. Sheth, B. Greenstein, D. Wetherall, G. Maganis, and T. Kohno. "Privacy Oracle: A System for Finding Application Leaks with Black Box Differential Testing," In CCS, 2008.

[15] D. Scott and R. Sharp, "Abstracting Application layer web Security," *In Proceedings of the Electrical International World Wide Web Conference (WWW 2004)*, Honolulu, HI, May 2002.

[16] NoScript.URL:
<https://addons.mozilla.org/en-US/firefox/addon/722/>

[17] Panopticlick.URL:
<https://panopticlick.eff.org/>

[18] Google Location. URL:
<http://www.google.com/support/websearch/bin/answer.py?answer=179386&hl=en>

[19] Trace My IP. URL:
<http://www.tracemyip.org>

[20] Browzar. URL:
<http://www.browzar.com>

[21] JAP. URL:
http://anon.inf.tu-dresden.de/index_en.html

[22] Safesurf. URL:
www.safesurf.com

[23] Wireshark URL:
www.wireshark.org

[24] Tcpxtract URL:
<http://tcpxtract.sourceforge.net>

[25] SQLite URL:
sqlitebrowser.sourceforge.net

[26] Facebook Profile ID URL:
http://www.ehow.com/how_5753004_facebook-id.html

[27] Network Miner URL:
networkminer.sourceforge.net

About the Author

Jason W. Clark is a security researcher who studies privacy and security issues of computer based systems. Mr. Clark has a Bachelor of Science in Information Technology from Syracuse University as well as a Master's of Science in Information Technology from Rensselaer Polytechnic Institute (RPI) and a second master's degree in computer forensics from George Mason University (GMU). Mr. Clark is currently working on his Ph.D. at GMU in the field of information technology/security with a focus on securing anonymizing network systems. He also works full-time at the Department of Defense as a lead information security analyst.