De-clawing Carnivore

CS588: Cryptology
Prof. David Evans

Dante Guanlao
Michael Tashbook
Leonard Woody
Dana Wortman

December 5, 2001

1 Introduction

> "They that give up essential liberty to obtain a little temporary safety deserve neither liberty nor safety." -- Benjamin Franklin, 1759.

The heated debate over Carnivore and online surveillance is reaching a fervid pitch after the September 11th events. Civil liberty groups claim that Carnivore's power is too far-reaching and encroaches on basic constitutional freedoms, while law enforcement agencies insist on needing more power to combat terrorism and crime. This report seeks to explain the debate over Carnivore and offer solutions for those wishing to elude the system as well as those who wish to make Carnivore a more viable option. The next sections detail Carnivore's technical implementation along with the legal context that surrounding it. Our report also offers solutions in which Internet users can evade Carnivore or combat its potential abuse by FBI agents. It is our hope that this information will be valuable to both sides of the debate as they struggle to find a balance between freedom and security.

2 Problem

2.1 Technical Aspects of Carnivore

Carnivore is a system of hardware and software combined to filter through internet traffic and find a designated target as outlined by a judge's written order. The system itself is not highly sophisticated. Understanding Carnivore will help give insight into the reasons civil libertarians are intensely concerned about the potential abuse of Carnivore's capabilities.

The entire system employs two or more computers and is comprised of components that could be bought in the public market today. One of the computers is located at the Internet Service Provider (ISP) of the target individual (referred to as the "collection computer"). The other computers involved are located at FBI centers (referred to as the "control computers"). These computers are used to process the intercepted data and manage the Carnivore software on the collection computer [Smith viii-ix]. Carnivore runs on Wintel machines with Windows NT/2000 as the operating system [Tyson]. The collection computer does not usually have a monitor or keyboard, but does have a removable Jaz disk drive [Smith viii-ix].

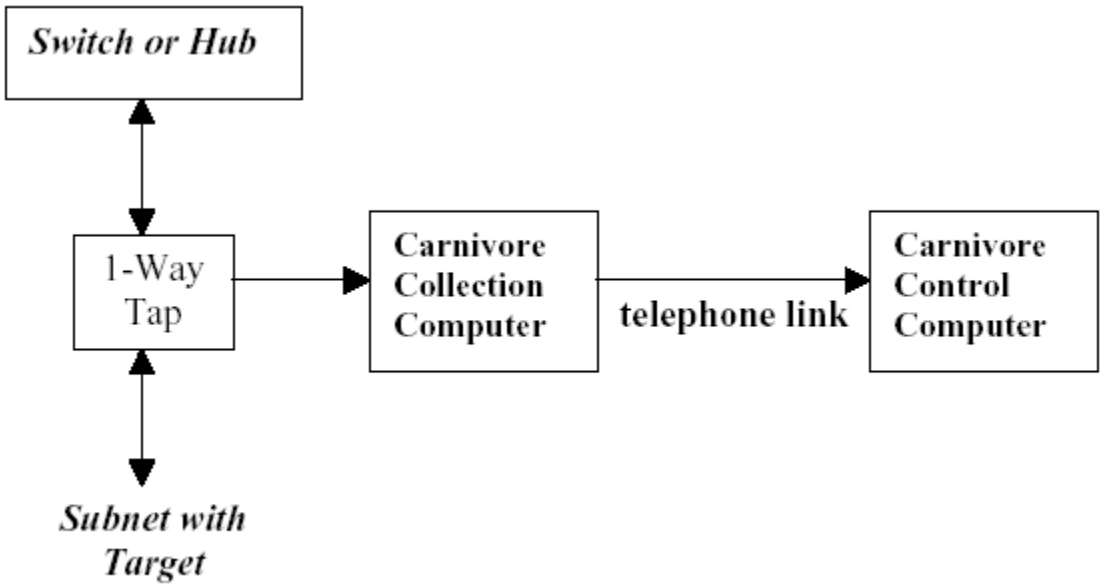

Figure 1 Carnivore Hardware System [Smith ix]
The diagram outlines the interconnections between the Carnivore hardware:

The Carnivore system is usually installed on a sub-network of the ISP that contains the targeted individual's internet traffic. The system uses a one-way tap into the traffic stream to collect its desired data. It takes each packet passing through the subnet and

sends a copy to the collection computer for filtering, and then sends the original packet to the hub or switch that normally processes the subnet packets. The collection computer also has a telephone link to the control computer(s). pcAnywhere, which is running on the collection computers, facilitates communication over the telephone link. These components comprise the communications hardware for Carnivore [Smith viii-ix, 3-10 – 3-13].

The communication between the control and collection computer is protected by two security systems: the encryption scheme of pcAnywhere and a challenge-type system. pcAnywhere is similar to SSH in that it uses a public key system to securely exchange a session key for symmetric encryption. The telephone link is further protected by a challenge-type system, in which a new challenge is generated each time the control computer attempts to connect to the collection computer. The challenge has to be answered correctly to connect. This is supposed to protect against an active eavesdropper. Unfortunately, there was not enough information to more adequately describe the challenge-type system employed [Smith 3-12 – 3-13].

What makes Carnivore different is its software. There are currently products on the market today that can look at packets going over the Internet (e.g. Etherpeek). But the FBI needed something that would comply with the narrow legal standards of wiretapping. Unfortunately, the FBI did not use a formal software development method to create Carnivore. Instead, they used a ″proof of concept″ method that produced many areas of concern that will be discussed later [Smith 4-6 – 4-7]. The software that composes Carnivore is made of four different programs. The first program is a Visual Basic GUI that enables agents to configure Carnivore. Second, a driver for interacting with the one-way tap has been written that evolved from a commercial product named WinDis 32. The third part is comprised of dll libraries that provide much of the functionality. The final part is an application program interface to use the WinDis 32 part of Carnivore [Smith ix, 3-18 – 3-19]. All these programs work together to give Carnivore its wiretapping capabilities.

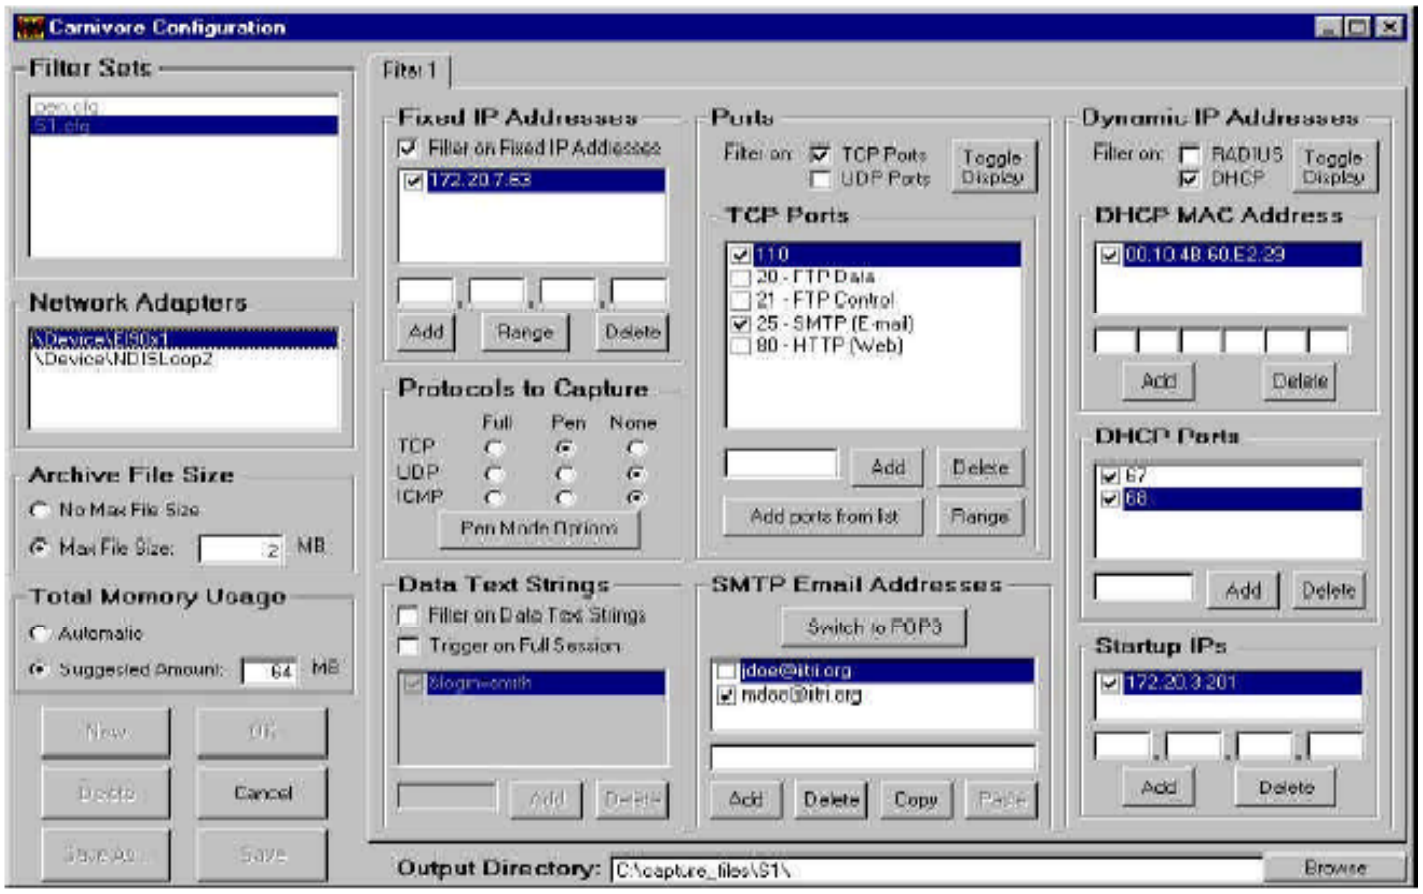

Figure 2 Carnivore Advanced Configuration GUI [Smith x]
Snapshot of the advance configuration GUI on the Carnivore collection computer

The main functionality of Carnivore is its filtering mechanisms. The GUI (pictured in Figure 2) is used to configure the filters to be used for the packets from the sub-network [Smith x]. There are myriad of ways that it can be setup. One type of filter is to save only packets that come from a fixed IP address. Carnivore is also designed to deal with dynamic IP addresses under the DHCP or RADIUS scheme. After selecting the IP address, an agent can choose under which protocols to collect packets: TCP, UDP, or ICMP. For each protocol, the agent can configure it as full, pen, or none mode. These correspond with the type of search orders issued and will be discussed in more detail later. If TCP or UDP protocols are selected, the agent can collect packets on all ports or just a chosen few. This facilitates listening to only certain types of communication (e.g. port 25 would give access to SMTP e-mail). Another type of filtering is to search for certain text strings within packets. This helps with web-based e-mail in searching for strings that are associated with particular information in the packets (e.g. ″to :″ and ″from :″). The final type of filter is to listen only to packets going to and from specific e-mail addresses under the SMTP and POP3 protocols [Smith 3-14 – 3-16]. However, it is very easy to take all packets for a particular protocol by simply choosing full mode for the protocol with no other filters set [Smith xi]. This is an extraneous function, since Carnivore should never be used this way under any type of judicial order. While this

presents a large variety of filtering options, it gives the agent a great deal of flexibility in setting up search orders, but some say the system is too flexible.

The pen and full modes of filtering are designed for two types of orders a judge can give and are tied to the history of wiretapping. When phones are tapped, the FBI can do a more limited collection termed a ″pen trap order″. These orders only allow the FBI to record the phone numbers called from a tapped phone and the source phone numbers of incoming calls. While in full wiretap orders, the FBI can collect specific communications from a certain individual or location [Smith 3-1 – 3-2]. Pen modes in Carnivore only take the e-mail packet headers or web packet IP addresses and replaces the remaining bytes of these packets with ″X″'s [Smith 3-15]. This has generated some controversy as it goes beyond the normal phone pen trap order, in that it gives the length of messages which could be volatile information.

This completes the technical description of Carnivore. It should be noted that Carnivore is the main part of collection of software called the DragonWare suite. The other two programs in DragonWare, Packeteer and CoolMiner, help with reconstructing and analyzing the collected packets [Smith viii-ix].

2.2 Legal Aspects of Carnivore

2.2.1 Introduction

As a tool for electronic surveillance, Carnivore falls under the same legal requirements and restrictions as traditional surveillance methods. In this section, we present a brief overview of the way in which surveillance is governed by existing laws

2.2.2 A brief history of wiretapping and United States judicial system

Assisted, undetected eavesdropping is a relatively new issue in the legal arena; prior to the development of long-distance communication technology like the telegraph and the telephone, it was not possible to eavesdrop without a greater risk of detection (consider, for example, the relative difficulty of unsealing and resealing a letter without leaving evidence). The subject immediately raises concerns over violation of the Fourth Amendment, which guarantees "[t]he right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures," and the Fifth Amendment, which protects individuals against self-incrimination.

The United States Supreme Court first considered the issue of wiretapping in 1928, in the case of *Olmstead v. United States,* 277 U.S. 438. Olmstead's appeal directly challenged the constitutionality of wiretapping; the Court, led by Chief Justice Taft, disagreed, ruling that wiretapping did not "amount to a search or seizure within the meaning of the Fourth Amendment." [Olmstead]

In the majority opinion for *Berger v. New York*, 388 U.S. 41 [Berger], Justice Clark provides an excellent summary of the Supreme Court's rulings regarding surveillance and the Fourth Amendment. Clark notes, in particular, that the Court in such cases has mainly focused on the issue of whether the surveillance required intrusion into a constitutionally protected area. Since Carnivore's monitoring equipment is located at the target's Internet Service Provider (ISP), this intrusion is not directly present, and the protection provided by the Fourth Amendment may not apply.

2.2.3 Procedural Issues Regarding Wiretapping

Wiretapping falls under the purview of two main pieces of legislation: Title III of the Omnibus Crime Control and Safe Streets Act of 1968, and the Foreign Intelligence Surveillance Act (FISA). [Smith] The passage of the Omnibus Crime Control Act was due partly to allegations that the FBI and CIA had conducted extended surveillance on civil rights leader Dr. Martin Luther King, Jr. in 1964. FISA governs the surveillance of foreign powers and their agents in the United States, and is concerned mainly with foreign intelligence rather than domestic criminal investigations.

Two types of wiretap are possible, as noted above. Pen register surveillance monitors outgoing calls from a particular communication device. Trap and trace surveillance monitors incoming calls. When the two are used together, as they often are, the result is known informally as pen-trap surveillance. [Smith]

Strict limitations are placed on electronic surveillance, in order to reduce the risk of abuse. Only certain high-ranking Department of Justice officials (the Attorney General and a few others) may authorize an application for a wiretap. Wiretaps can only be ordered for the investigation of felony offenses, and only upon a showing both of probable cause and that normal investigative measures are not sufficient. In addition, the scope of the surveillance is restricted to some specific criminal activity; investigators are only permitted to eavesdrop on those communications that are related to the particular crime being investigated. This requirement, known as *minimization*, is intended to limit the intrusion on the target's Fourth Amendment rights. Finally, wiretap orders can only be granted by a select set of judges and state courts—specifically, those established under Article III of the United States Constitution. Article III judges are protected from political and financial pressure (they are not subject to arbitrary removal, and their salaries cannot be reduced), which further reduces the possibility that wiretaps may be ordered for the wrong reasons (political vendettas, for example). [Smith]

18 USC §3123 sets out a series of requirements that must be satisfied before a pen register or trap and trace order can be issued. An application for surveillance under this section must specify:

1. the identity of the owner of the telephone line that will be tapped,
2. the identity of the individual who is the subject of the criminal investigation,
3. the number and physical location of the telephone line that will be tapped (for a trap and trace device, the application must also specify the geographic limits of the trap and trace order), and

4. the offense for which relevant information is expected to be collected.

Surveillance orders granted under this section may not exceed sixty days in duration, although agents may apply for an additional extension of sixty days. [U.S. Code] This restriction also supports the minimization requirement. Extended surveillance is not permitted, so as to reduce the intrusion incurred by the wiretap.

Internet traffic poses a problem for Carnivore in terms of the minimization requirement. TCP/IP packets contain content along with addressing information. This means that it is possible for Carnivore to exceed the scope of a wiretap order by examining the content of packets, instead of examining addressing information only.

In 1986, the wiretapping laws were modified to allow "roving" wiretaps. [ACLU] Prior to this amendment, it was possible to thwart a wiretap simply by using other telephones (since a wiretap order was bound to a specific telephone number). If law enforcement personnel could show that the target of their surveillance was deliberately using different telephones in order to escape their wiretap, they could seek to obtain a "roving" wiretap that would apply to an individual rather than a specific phone number.

The standard for granting a roving wiretap was relaxed in 1998. Under this looser standard, a roving wiretap could be obtained simply for the reason that the target used multiple telephones. It was no longer necessary to show that the target was deliberately attempting to evade surveillance. [ACLU] Any phone near the target can be tapped under a roving wiretap, on the assumption that it might be used for some criminal purpose. In an attempt to prevent eavesdropping on innocent persons, the wiretap statute was modified to require law enforcement personnel to ascertain that the target was actually using a particular phone line before tapping it. The ACLU notes that the Supreme Court has yet to rule on the constitutionality of roving wiretaps; they may yet be determined to violate the Fourth Amendment search and seizure guarantee. [ACLU]

How do these restrictions interact with the evasion tactics described elsewhere in this paper? Crowds increase the chances that Internet traffic from a given machine may actually have originated somewhere else. As a result, it is more difficult for law enforcement personnel to track a target without tapping a large number of machines. In the process, they run the risk of violating the privacy of a large number of innocent people, since there is no way to tell ahead of time whether a given message will be routed through a particular machine. A similar situation arises for onion routing; onion routing makes it necessary to place wiretaps on a large number of machines (possible onion routers), on the off chance that a relevant message will appear.

2.2.4 Recent Developments

The privacy landscape changed dramatically after the terrorist attacks of September 11, 2001. In response to the terrorist attacks of that date, a new bill, dubbed the "USA

PATRIOT[1] ACT of 2001," was introduced into the House of Representatives. The Patriot Act was signed into law on October 26, 2001, and arms law-enforcement personnel with greatly expanded wiretapping powers. In particular, the Patriot Act removes several of the procedural checks and balances that had previously limited the government's ability to conduct electronic surveillance.

In his dissenting opinion in *Olmstead*, Justice Brandeis wrote:

> Experience should teach us to be most on our guard to protect liberty when the government's purposes are beneficent. Men born to freedom are naturally alert to repel invasion of their liberty by evil-minded rulers. The greatest dangers to liberty lurk in insidious encroachment by men of zeal, well-meaning but without understanding. [Olmstead]

This observation is one cited by many privacy advocates who object to the passage of the Patriot Act in the wake of September 11th. Privacy advocates have expressed concern over the rapid deployment of the bill; an analysis by the Electronic Frontier Foundation [EFF] observes that key procedural processes, such as inter-agency review and the normal committee and hearing processes, were suspended for the Patriot Act when it was introduced. The EFF also points out that the authors of the Patriot Act failed to present compelling evidence that the (then) current level of government surveillance powers were inadequate for the purposes of investigating and prosecuting acts of terrorism. [EFF] Several hours after the Patriot Act was signed into law, the FBI immediately expanded its electronic surveillance activities, taking advantage of its new powers to monitor cable modem users without judicial approval. [CNN, WIRED]

One significant change to the wiretap laws under the Patriot Act reduces the level of proof required to obtain a wiretap. [Patriot Act, section 216] Previously, applicants were required to provide probable cause; now, they need only assert that the information to be obtained is "relevant to an ongoing criminal investigation." The ACLU notes that, with this change, judges have far less ability to protect individual privacy rights; wiretap orders must be granted, even when there is very little evidence that the results will in fact be relevant to the investigation in question. [ACLU]

The Patriot Act also greatly extends the reach of electronic surveillance efforts. Previously, wiretap orders were restricted to a specific jurisdiction (that of the court issuing the order). Now, however, "terrorist investigations" have been added to a select list of crimes with a single, national jurisdiction. This means that wiretaps can be used nationwide with a single order; even if a target uses multiple ISPs in different parts of the country, law enforcement personnel need only obtain a single wiretap order in order to eavesdrop on any of those connections. [Patriot Act, section 219]

[1] Here, "USA PATRIOT" is an acronym for "Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism."

Finally, section 224 of the Patriot Act outlines a "sunset provision," which specifies that its amendments to existing statutes will expire on December 31, 2005. A number of changes are exempted from this sunset provision, however, including several critical modifications. In particular, section 216, which reduces the amount of proof required for a wiretap application, and section 219, which allows single-jurisdiction search warrants to be issued for terrorist investigations, do not sunset. [EFF, Patriot Act] At present, it is not clear whether privacy advocates will succeed in their efforts to overturn these modifications.

3 Related Work

The primary work that describes Carnivore and was used as a foundation for this report was the IIT Research Institute's (IITRI) *Independent Review of the Carnivore System*. The Department of Justice (DoJ), under pressure from civil liberty groups and Congress, asked for proposals to be submitted for evaluating Carnivore. IITRI was contracted to be the evaluator. The DoJ requested four specific criteria to be investigated about Carnivore's functionality:

1. Provide [FBI] investigators with all, but only, the information it is designed and set to provide in accordance with a given court order.
2. Introduce any new, material risks of operational or security impairment of an ISP's network.
3. Risk unauthorized acquisition, whether intentional or unintentional, of electronic communication information by:
   (1) FBI personnel
   (2) Persons other than FBI personnel
4. Provide protections, including audit functions and operational procedures or practices, commensurate with the levels of the risks. [Smith]

Criteria 1, 3, and 4 directly dealt with our project. IITRI was given the source code for Carnivore and performed extensive testing of Carnivore. The report has over 100 pages on the details and functionality of Carnivore. Therefore, we thought it was a resource of supreme value in carrying out our project.

The report had many conclusions. Many of them were critiques of the current Carnivore implementation. We picked conclusions out and expounded upon them and found critiques of our own. For example, the report concluded that significant potential existed for abuse of the system by FBI agents [Smith]. We specifically developed a way to combat that. We also surveyed ways to evade Carnivore. The *Independent Review of the Carnivore System* was an excellent starting point for our project.

In order to understand what strategies for anonymous and encrypted internet communication were applicable to defeating Carnivore, we looked at many different ideas for both real-time (such as HTTP) and non-real-time (such as SMTP) communication. The most resourceful documents are as follows:

Naval Research Laboratory. "Onion Routing." – The website that documents the Onion Routing system as developed and tested by the Naval Research Laboratory. Along with a substantial amount of technical information, it also provides the user with practical timings and examples of the system in action. Most importantly it gave a general overview of onion routing.

Goldschlag, David, Michael Reed, and Paul Syverson. "Onion Routing for Anonymous and Private Internet Connections." – Article that describes Onion Routing at a high level, but also discusses some of the future possibilities of onion routing, such as combining it with an anonymous remailer to protect the server and receiver's identities. It also covers some of the potential drawbacks to the onion routing system, such as computational overhead, but then excuses those as unnoticeable compared to other internet traffic delays.

"The International PGP Homepage." – This site describes the PGP program including the history of PGP, specification and standards documentation, and an explanation of how PGP works. Understanding how PGP works helps one to understand how Carnivore is defeated by it and why so many security schemes depend on the unbreakable aspects of PGP encryption.

Reiter, Mike and Aviel Rubin. "Anonymous Web transactions with Crowds." – Article that provides a basic understanding of crowds including defining what level of security is afforded by crowds. Although concentrated on web transactions, the policies developed in crowds can be combined with some of the other available schemes and can provide both real-time and non-real-time anonymity as well as security.

Reiter, Mike and Avi Rubin. "Crowds: Anonymity Loves Company." – Website associated with the "Anonymous Web transactions with Crowds" article. Essentially presents similar information, but develops the primary risks of crowds in an easy to understand section.

Hetrick, Brian. "Personal Security: Pseudonymity." – Website that describes nym servers: what they do and how to use them. It is essential to understand how this most basic concept can defeat Carnivore, but also how susceptible it is to other forms of attacks such as timing attacks and traces.

Mazières, David and M. Frans Kaashoek. "The Design, Implementation and Operation of an Email Pseudonym Server." – Technical document detailing a particular implementation of a nym server and the problems that were faced and, for the most part, successfully dealt with. Understanding the limitations of a nym server helps us to understand better how Carnivore might be able to successfully connect an anonymous sender to a particular message.

Pfitzmann, Andreas, Birgit Pfitzmann, and Michael Waidner. "ISDN-MIXes: Untraceable Communication with Very Small Bandwidth Overhead." – Document that describes an implementation of mixes designed to work on a bandwidth limited system (ISDN). Mixes allow both the sender and receiver to remain anonymous with regard to each other as well as the message path itself to remain untraceable.

Berthold, Oliver, Hannes Federrath, and Marit Köhntopp. "Project 'Anonymity and Unobservability in the Internet'." - Study (and comparison) of several major schemes for anonymous or untraceable communication. Discusses the pros and cons of

each and then combines the best features of these into a revised mix. This is also an Excellent source for understanding the drawbacks for each strategy.

4 Solution

4.1 Strategies for Avoiding Carnivore

There are a variety of strategies that can be employed to defeat the Carnivore system. Some of these strategies focus on providing anonymity to the sender and receiver of an email or for the sender from the receiver of an HTTP request. Other strategies focus on encryption schemes that can be employed to hide the contents of the message from an eavesdropper. Several incorporate both ideas into a successful blend of anonymity and intractability. We will describe several of these strategies and how they work to defeat or evade the Carnivore system.

4.1.1 PGP

PGP (Pretty Good Privacy) is a program that employs varied cryptographic techniques to establish a method for secure communication among individuals, businesses, and other entities. PGP works by first establishing a public-private key pair for an entity. The public portion is published using a trusted third party and the private key portion is kept secret by the entity. For each message the entity wishes to encrypt, PGP then creates a key to be used only for that message (session key), compresses the message to be sent, and then encrypts it using a fast conventional encryption algorithm with the session key. PGP then encrypts the session key using the public-key encryption algorithm, which is significantly slower than the conventional algorithm, and attaches this to the encrypted message.

PGP eludes Carnivore by employing encryption to disguise the text in the message. Unfortunately, since the original receiver and sender are known from the message header which must remain in plaintext, Carnivore can still collect the encrypted packets from an individual who is under surveillance. Later, an FBI agent can use equipment to attempt to break the encryption. Fortunately, it is believed that no one has the ability to break public-key encryption at this time and so encrypted information is relatively secure.

4.1.2 Nyms

Nyms are anonymous remailer servers where the sender's identity is hidden from the recipient. A user can configure the nym server by giving it a set of instructions on how to send a reply message through a set of anonymous remailers. When someone sends a message through the nym server, they first send a message through a set of anonymous remailers and when it reaches the nym server, the server rewraps the message and sends it to the intended recipient including a reply-to email that is associated with the nym server. Upon receiving a reply from the message, the nym server follows the instructions for returning the reply through the anonymous remailers. In addition, the server itself does

not maintain logs or information relating the email address to the sender with the exception of the instructions for submitting a reply to an anonymous remailer.

Two of the key features of a nym server are that all message traffic is encrypted and any single hop only knows the previous and next remailers in the chain. Therefore, without concerted effort at tracing the encrypted messages through the remailers, or without coercing each node to reveal the next node in the series, it is virtually impossible to determine the sender or receiver. Both the encryption and the anonymizing nature of the message transfers makes it extraordinarily difficult for Carnivore to determine which messages or replies belong to the observed sender or what those messages might contain.

4.1.3 Crowds

Crowds is an anonymous web-browsing technique that allows a user to make HTTP requests without the receiving server knowing the sender's identity. Anonymous web-browsing is accomplished through the use of a specialized proxy (jondo) that randomly chooses to forward the HTTP request to another member of the crowd or to send it to the intended server. Communication between jondo's is encrypted, thus protecting from a passive eavesdropper within the crowd. However, any jondo can read the information in the message as can anyone eavesdropping over the link between the final jondo and the intended server which is not encrypted. The receiving server has no way of telling which member originated the request as it is just as likely to have originated from any member of the crowd. It is also impossible for the individual members of the crowd to know which jondo originated the request as the original sender is indistinguishable from a jondo that just forwarded the request. Mathematically, anonymity can be maintained with anything less than 1/3 of the jondo's collaborating. It is worthwhile to note that while Crowds, with a very high probability, promises anonymity, it makes no assurances about the privacy of the data in the message.

Crowds serves as a stumbling block for Carnivore in that Carnivore cannot associate web-traffic from a particular machine as being traffic from the user of that machine. Any messages that Carnivore intercepts from that machine are just as likely to be from any other jondo in the Crowd. One of the weaknesses in Crowds is since only the communication internal to the Crowd is encrypted, any message that is intercepted between the Crowd and the intended server will not be encrypted making the contents of that message readily accessible to Carnivore.

4.1.4 MIXes

MIXes is a secure routing scheme in that the sender of a message encrypts the message with the keys of each node along the path. An unique point being that the sender can remain anonymous to the receiver as well as the receiver remaining anonymous to the sender. The system works such that the sender and receiver both send messages to a

"middle-man" who will connect the two streams of data with a label that was broadcast anonymously from the sender earlier. When the middle-man receives an encrypted message from the first string of routers (MIXes), it decrypts the message, interprets the label, and then passes it onto another string of MIXes to the correct reciever. The original sender encrypts the message with the key of each MIX on the way to the middle-man (including the middle-man's key). Each MIX as it receives a message, decrypts it and then passes it along to the next MIX. On the receiving side, the middle-man simply encrypts the message once and then passes it along to the first MIX. As each MIX receives the message, the MIX encrypts the message with its own key and passes it along. So, the receiver gets a message that has been encrypted with the keys from each of the MIXes along the return path from the middle-man.

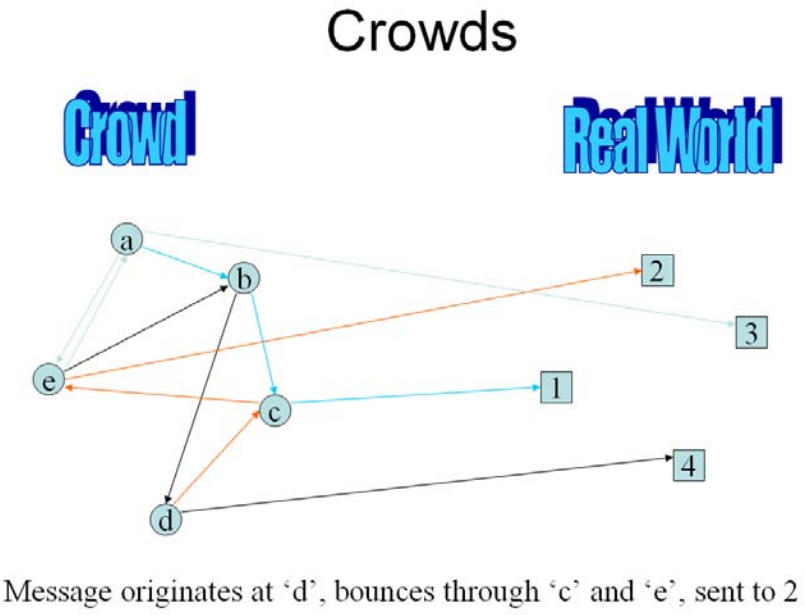

Figure 3 Crowds: Traffic Description
A depiction of the message traffic through a crowd and then transmitted to the intended receivers in the real world.

From the outside, all message transmissions are encrypted, but in addition, each node knows only the identity of the previous node and the next node, so without the cooperation of all the nodes, Carnivore would not be able to connect a sender or receiver with a particular message. MIXes also work against timing-related attacks in that each MIX collects several messages, mixes up their order, and then sends them along to the next MIX. This can defeat the FBI's attempt (external to Carnivore) to connect sending and receiving patterns through the MIX system.

4.1.5 Onion Routing

Onion Routing is a communication technique based on MIXes that supports anonymous connections, but not necessarily anonymous communication. Users of an onion routing system may expect to have privacy from eavesdroppers such that the eavesdropper cannot

determine who is talking to whom, or the contents of any transmission. An application that wishes to use an onion routing system needs only to connect to an onion router proxy and the proxy then creates a list of routers to send the message through. For each proxy on the route, the initial proxy concatenates to the message a header specifying the next proxy in the chain and then encrypts the concatenation with the proxy's key. As each proxy on the route receives the message, the proxy first decrypts the message using the key, then, if the message is not intended for its machine, the proxy then forwards the message to the next proxy in the chain as defined by the header. If the proxy is the intended receiver, it then passes the message to the receiving application.

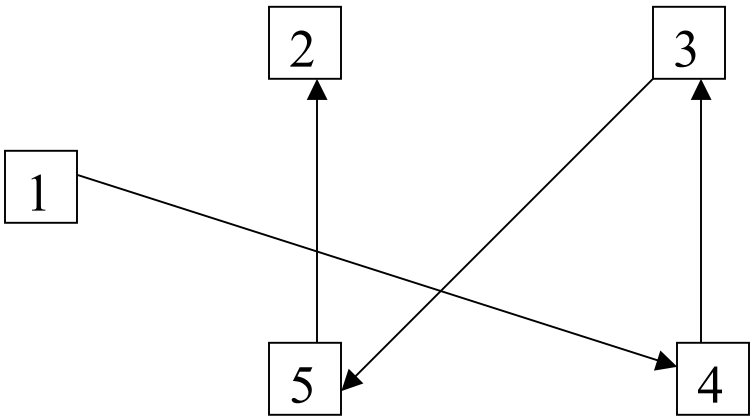

Message = E_{k4}("send to 3" | E_{k3}("send to 5" | E_{k5}("send to 2" | E_{k2}(M))))

Figure 4 Depiction of an onion routed message.
Each square represents an onion routing proxy. The original sender (1) sends the message through hops (4), (3), and (5) in order to reach (2).

Onion Routing can be used to avoid Carnivore in several ways. First, onion routing uses encryption, which Carnivore alone cannot break. Therefore, when Carnivore collects a packet, it can only search the header information for matches with the collection criteria, it cannot search the body of the message since it is encrypted. Second, since any particular communication between proxies only has the identity information for a single hop in the chain, anyone who collects this information will not know the original sender or the final receiver; therefore, the original sender and final receiver remain anonymous to someone eavesdropping along any hop.

4.2 Strategies for Strengthening Carnivore

The Carnivore System, in its current implementation has many weaknesses that are a direct result of its "proof of concept" creation as discussed earlier. This section will

analyze some of the greater weaknesses and recommend ways in which the system may be better implemented.

The largest failing in the Carnivore System is its vulnerability to rogue FBI agents abusing the system. As discussed earlier, the choice of one radio button in the Carnivore GUI will change the system from a simple Internet wiretapping device to a full-fledged monitoring system for anyone using the ISP. A more secure implementation of Carnivore would include an encryption system that would minimize the amount of flexibility that individual agents have in configuring the system. For instance, a judge would issue, along with the court order to use Carnivore on an ISP, an encrypted message that, once put into Carnivore would configure the system to do only what the courts have allowed. A full explanation of the encryption trust model follows:

1. A judge issues a court order. Along with this court order is a message encrypted with the judge's private key which includes multiple hashed messages corresponding to the different configuration choices on the system, as well as a hash of his public key for identification verification purposes somewhere in that list of hashed Carnivore settings.

   $$\mathbf{E_{KRA}[H(M_1), H(M_2), \ldots, H(M_{K-1}), H(E_{KUA}), H(M_K), H(M_{K+1}), \ldots, H(M_N)]}$$

2. The FBI agents will then take that message, input it into Carnivore, and notify the system that the following message is from Judge X
3. Carnivore will then connect to an FBI database of judges' public keys over a secure channel and pull Judge X's key from the database, decrypt the message and then hash the gotten public key, comparing the result to the list of hashed messages.
4. Upon finding a match to the hashed public key gotten from the database to the public key hashed in the original encrypted message, the system will then set itself to the configurations as defined by the hashed messages and then proceed as the current implementation of Carnivore is defined.

This encryption scheme has many advantages over the current implementation of Carnivore, but also makes some assumptions. Some of the advantages are:

1. Since the system is being set by electronic code and not individual agents, the probability for rogue agents tampering with the system is reduced.
2. The use of the judge's public key both as a hashed part of the message and to decrypt the message serves as double verification of who the judge says they really are. Should another judge's public key be gotten from the public key database (possibly caused by a rogue agent telling Carnivore that the encrypted message came from Judge Y when in reality it came from Judge X) and decrypt the message, a hash of that other Judge's public key will not result in a match anywhere in the message and cause the system to halt. In addition to this, for a judge to issue a court order for this encryption scheme, they must first be verified by the FBI to hold their public keys in their database.

3. Should a rogue agent get the judge's public key and decrypt the message issued by the judge, they will still have to guess at which part of the resulting message is the hash of the public key, and what parts correspond to the correct control settings being hashed.
4. This scheme allows for partial accountability since only one judge has the specific key for configuring Carnivore.

However, this trust model does rely on some assumptions:

1. The hash function contained within Carnivore must satisfy the properties of a good hash function.
2. There must be a trustworthy judge that will set the control messages correctly for Carnivore.
3. To get the public key for verification, there must be a secure channel between the Carnivore system and the FBI database. Without this, an active eavesdropper could intercept the public key request and then send a false public key, halting Carnivore.
4. Since the development and implementation of Carnivore will be done completely by the FBI, the FBI must ensure that its designers correctly implement the Carnivore system through a formal software development method.

The above trust model deals with the possibility of rogue agents. However, there are still failings to the Carnivore system. The next paragraphs will take those issues discussed earlier and recommend ways in which they may be fixed.

1. **Accountability.** The Carnivore system as it stands does not log who the agents are in charge of running the system, or the judge who issues the court order. Expanding upon the partial accountability as described in the encryption scheme will help keep track of who is involved should something go wrong with using Carnivore.
2. **Formal Software Development.** The "proof of concept" method for building Carnivore leaves gaps in knowledge and documentation system. A formal development method will more clearly explain what Carnivore is intended to do in given situations as well as provide accurate documentation for troubleshooting purposes.
3. **Encrypted Message Length**. The pen trap method and its replacement of the message with X's can still provide the reader with knowledge of the length of the message. Further implementations should either pad the X's to some predefined length (e.g. 2000 characters) or eliminate them completely.

In any large software project, there are an expected number of flaws and bugs created in the process. This part of the report has hopefully addressed these concerns and by following these recommendations, the FBI will be able to safely and securely use Carnivore. These recommendations will also reduce the number of vulnerable points where Carnivore could be exploited.

5 Evaluation

5.1 Evaluation of strategies for avoiding carnivore

The strategies for defeating Carnivore, essentially utilize two primary methods for evading Carnivore: confusion of the delivery path and encryption of the message. Confusion of the delivery path primarily results in anonymity to the sender and/or receiver. Encryption of the message simply makes it difficult to determine the contents of the message, but does not provide any anonymity to the sender or receiver. Those strategies that utilize both methods for the entire journey of a message are the most successful. According to Berthold, et. al., there are six primary attacks that the perfect system must be able to withstand (paraphrased from Berthold):

1. Message Coding Attack – a message that does not change encoding throughout transmission can be tracked
2. Timing Attack – watching presumed start and end points of a communication and observing correlations between the beginning and the ending of connections
3. Message Volume Attack – a message that does not change length significantly can be tracked
4. Flooding Attack – an attacker can flood a system with messages to eliminate all other messages but the one he or she wishes to observe
5. Intersection Attack – an attacker can observe a user over a long period of time and can correlate messages with logged-on and logged-off periods of time
6. Collusion Attack – a coalition of users or systems can track a user through the system

Anonymous remailing systems provides little to no protection against any of these attacks, except when used in conjunction with encryption. With encryption, it provides some protection against message coding attacks by encrypting the message between proxies, therefore changing the substance of a message during transmission. Nym servers provide protection against several of these attacks and with the additions suggested by Mazières and Kaashoek (Mazières), the server can stand up to almost all of these attacks to some degree. Crowds offer little protection against anything except the message coding attack as the message is encrypted throughout the transmission through the crowd. However, the message is not encrypted once it is sent to the real world, so there is no protection against the contents being collected, even though an outsider cannot trace it back to the sender. Onion routing offers quite a bit of protection within the routing network, but does not provide protection at the end points. Mixes provide an additional layer of protection at the endpoints, but at a very high cost since there are many iterations of encryption.

Unfortunately, none of the above strategies is perfect with the intersection attack being the most likely attack to withstand all current strategies as there are currently no methods for defeating this attack. Since even the weakest of these strategies is capable of evading Carnivore, it can be observed that Carnivore is an extraordinarily simple system in that it takes few active steps in attempting to determine the sender of a message, or track a

message from sender to receiver. Carnivore has been created such that it uses the easiest way to try and find the sender of message and therefore can easily be defeated by employing a strategy that could collapse under the simplest of attacks. Therefore, it seems as though it would not be difficult to develop a system similar to Carnivore that would have the capabilities of employing some, if not all the of the above attacks in attempting to determine the actual sender of a message that has been sent through a proxy-based system. Unfortunately, working under the assumption that strong encryption has not been broken by the FBI, the contents of those messages would not be able to be read under many of these schemes, but at least it could be determined that a particular individual has been communicating with another individual.

5.2 Evaluation of strategies for strengthening carnivore

The combination of adding these features to Carnivore will make it more secure both for those using it and for those under surveillance. However, simply adding features to a flawed system will not fix everything. To better implement Carnivore, it must be rebuilt from the ground-up using the Formal Software Development Method as described before. Only then could some guarantee of safety along with security be issued.

6 Conclusion

It should be clear now why the debate over Carnivore is so heated. The previously described gaps and loopholes explain why the federal government feels the need to update and improve Carnivore for better combating terrorism and crime. The legal standards give insight into the justification civil liberty groups use when they argue against Carnivore. We hope our solution for strengthening Carnivore will help the system to be used legally and efficiently. Our description of evasion techniques should show civil liberty groups the real danger that exists with criminals who are tech-savvy. This report not only helps the experts of law enforcement, but also educates the general public about this important debate.

7 Bibliography

(ACLU) American Civil Liberties Union. "How the USA-PATRIOT Act Limits Judicial Oversight of Telephone and Internet Surveillance". October 23, 2001. Available online at <http://www.aclu.org/congress/l102301g.html>

(Berger) Berger v. New York. Cited as 388 U.S. 41 (1967). Available online at <http://laws.findlaw.com/us/388/41.html>

(Berthold) Berthold, Oliver, Hannes Federrath, and Marit Köhntopp. "Project 'Anonymity and Unobservability in the Internet'." Computers, Freedom and Privacy. Proceedings of the tenth conference on Computers, freedom and privacy : challenging the assumptions, 2000. ACM Press New York, NY, USA. Pages: 57-65.

(CNN) "Bush Signs Antiterrorism Bill Into Law". CNN.com, October 26, 2001. Available online at <http://www.cnn.com/2001/US/10/26/rec.bush.antiterror.bill/index.html>

(EFF) Electronic Frontier Foundation. "EFF Analysis of the Provisions of the USA PATRIOT Act That Relate to Online Activities". October 31, 2001. Available online at <http://www.eff.org/Privacy/Surveillance/Terrorism_militias/20011031_eff_usa_patriot_analysis.html>

(Goldschlag) Goldschlag, David, Michael Reed, and Paul Syverson. "Onion Routing for Anonymous and Private Internet Connections." Communications of the ACM, Volume 42, Issue 2 (February 1999). Pages: 39-41.

(Hetrick) Hetrick, Brian. "Personal Security: Pseudonymity." np. Online. Internet. 15 November 2001. Available: <http://www.geocities.com/tnotary/sppnyms.html>

(International) "The International PGP Homepage." np. Online. Internet. 20 November 2001. Available: <http://www.pgpi.org/>

(Mazières) Mazières, David and M. Frans Kaashoek. "The Design, Implementation and Operation of an Email Pseudonym Server." Conference on Computer and Communications Security. Proceedings of the 5th ACM conference on Computer and communications security, 1998. ACM Press New York, NY, USA. 27-36.

(Naval) Naval Research Laboratory. "Onion Routing." np. Online. Internet. 20 November 2001. Available: <http://www.onion-router.net/>

(Olmstead) Olmstead v. United States. Cited as 277 U.S. 438 (1928). Available online at <http://laws.findlaw.com/us/277/438.html>

(Patriot Act) USA PATRIOT ACT. Full text as of October 1, 2001. Available online at <http://www.politechbot.com/docs/patriot.act.100101.pdf>

(Pfitzmann) Pfitzmann, Andreas, Birgit Pfitzmann, and Michael Waidner. "ISDN-MIXes: Untraceable Communication with Very Small Bandwidth Overhead". Proceedings of GI/ITG-Conference "Communication in Distributed Systems). 20-22 February, 1991. Pages: 451-463.

(Reiter 1) Reiter, Mike and Avi Rubin. "Crowds: Anonymity Loves Company." AT&T Research Labs. np. Online. Internet. 23 November 2001. Available: <http://www.research.att.com/projects/crowds/>

(Reiter 2) Reiter, Mike and Aviel Rubin. "Anonymous Web transactions with Crowds." Communications of the ACM. Volume 42 , Issue 2 (February 1999). Pages: 32 – 48.

(Smith) Smith, Stephen P., et. al. "Independent Review of the Carnivore System: Final Report". IIT Research Institute. 28 October 2001. <http://www.usdoj.gov/jmd/publications/carniv_final.pdf>

(Tyson) Tyson, Jeff. "How Carnivore Works". 28 October 2001. <http://www.howstuffworks.com/carnivore.htm>

(U.S. Code) United States Code. Available online at <http://www4.law.cornell.edu/uscode/>

(WIRED) McCullagh, D., and Polen, B. DOJ's Already Monitoring Modems. WIRED magazine, November 28, 2001. Available online at <http://www.wired.com/news/conflict/0,2100,48711,00.html>