Freedom 2.1 Security Issues and Analysis

Adam Back, Ian Goldberg, Adam Shostack
Zero-Knowledge Systems, Inc.
{adamb,ian,adam}@zeroknowledge.com

May 3, 2001

Abstract

We describe attacks to which Freedom, or Freedom users, may be vulnerable. These attacks are those that reduce the privacy of a Freedom user, through exploiting cryptographic, design or implementation issues. We include issues which may not be Freedom security issues which arise when the system is not properly used. This disclosure includes all known design or implementation flaws, as well as places where various trade-offs made while creating the system have privacy implications. We also discuss cryptographic points that are needed for a complete understanding of how Freedom works, including ones we don't believe can be used to reduce anyone's privacy.

1 Introduction

Readers not regularly exposed to security work may not know that the publication of analysis is an important part of how security professionals work. Open discussion is the best way we know to improve the security of systems we create. Please understand as you read this paper that there is no such thing as perfect security. One well known expert, Bruce Schneier, has said "The only secure computer is one that is turned off, locked in a safe, and buried twenty feet down in a secret location—and I'm not completely confident of that one, either." We choose to disclose all the known security issues that Freedom has because we believe that this is the right thing to do.

We don't mean to scare anyone away from the system; we believe it offers solid protection against many threats, and is better than the alternatives. We note with disappointment that no other privacy company has chosen to publish such a document. Our intent is to constantly improve the system and make it better and we will make available a new security analysis with each version. Many of the attacks here do not apply to competitive systems, not because they are immune, but because they fall to simpler attacks.

We have used a number of methods to find problems. We did a roundup of the product's programmers to determine what they know to be broken or what has been worrying them. All programmers also had a chance to review all of our white papers, and the most experienced have reviewed them in-depth for comments and to find more discrepancies.

The cryptography in the system has been reviewed by a number of experts, who have collectively pointed out a large number of flaws and issues, most of which have been corrected, and as such are not noted here. Those flaws which remain are our fault, not theirs.

In addition, a few experienced consultants have done walkthroughs and code reviews, and the problems that they found were corrected.

2 Statement of Security

The Freedom system delivers the highest quality privacy protection available to the consumer today. Freedom has been designed to protect the privacy of users sending email, browsing the web, posting to news groups and participating in Internet chat. Freedom's privacy protection is designed so that even if Zero-Knowledge wanted to violate your privacy, we'd have trouble doing it. Freedom was also designed to ensure that none of our partners can violate your privacy, intentionally or accidentally.

Freedom is not invulnerable — no system is. We've done our best to make it very, very difficult, time-consuming, and expensive for an attacker to break. One of our goals is to offer the best protection available to the consumer today, and we believe that we have succeeded. There are, however, several known technical issues which may lead to breaches of privacy, and the purpose of this paper is to share that knowledge with you, because that is the right thing to do.

We estimate that it would require the resources and dedication of a large or dedicated intelligence agency or a dedicated internal surveillance force to effectively, reliably, and on an ongoing basis, break the privacy which we offer with regard to web browsing, chat, or other interactive services.

One way to attack a user would be to send them a large mail, and then when they connect via the freedom network to the mail system to pick up their mail, to try and track back the connection through the network to the user. The route is randomly selected on each connection, and only lasts for 1/2 an hour. It is possible that a very powerful attacker could compromise machines or record traffic at enough network access points to identity the user.

Zero-Knowledge considers this better than available alternatives, all of which except Cypherpunk and Mixmaster remailers[1] can be compelled to compromise your privacy with a single warrant.

A backbone provider may be able to monitor many links, possibly gathering much data on its own behalf, or more likely, in collusion with a law enforcement agency.

A hacker group may be able to engage in attacks that approach the abilities of a national intelligence agency; however, it seems unlikely that they can translate the compromise of the targeted networks into an ongoing intelligence gathering operation with data gathering, storage, analysis, summation and dissemination. It is much more likely that they can verify a guess as to the identity of a nym, or engage in a set of targeted compromises to discover what nym a targeted user has or is using.

[1] With remailer reply blocks, the user can choose a number of hops, and their location.

Lastly, we note that the Freedom system is vulnerable to denial of service attacks. We do not enumerate these here because we don't consider most of them to be interesting, but rather, annoyances exploited by the immature.

3 Overview of Threats

In making claims about the protection that we offer, and in explaining the limits of our service, it is useful to examine some of the types of people who may attempt to violate your privacy. Below, we briefly describe those attackers and our assumptions about their abilities.

3.1 Web Site Operators

A web site operator can offer cookies, and send you 'active content' to try to track you. Many web sites will use various forms of encouragement to get personal information about you, such as asking for your ZIP code for weather reports, and then share that information with their advertising networks. The advertising network, by placing ads on many sites, is able to gather a large profile of you. Internet sites using custom protocols, like Real, can also engage in tracking of users.

Web sites can also use ActiveX, Javascript, and other languages to cause your computer to send information to the site. This behavior is more unusual than gathering profiles through cookies.

3.2 Sysadmins

Systems administrators can variously read your mail, watch where you make network connections (such as web browsing), and generally monitor all your unencrypted online activities. Your company sysadmin can read any files you store on network drives, and may also be able to access all the files on your desktop or laptop computer. There may be laws in your area controlling this activity, and you may have signed away all of your rights under such laws as part of an employment contract.

3.3 Search Engines

Search engines can discover an awful lot of information that you, your friends and family, your employer, your school or alma mater, and others in your life may have placed online.

3.4 Lawmakers and Law Enforcement

In democracies or other countries where the police are under the jurisdiction of civilian authorities, police or court threats are usually overt, in the form of attempts to obtain encryption keys to force data recovery, including identity information. This is usually involves warrants or court orders, but may also include psychological tactics or even physical intimidation.

In some countries, police may also operate covertly through actions such as emissions monitoring and "dumpster diving." One cannot assume that all police actions are authorized or even legal, or that if authorized and legal, the regime that has authorized them is ethical and protective of human rights. Police in various countries have been known to use illegal means of gathering information, which they abandon when it leads them to a legal way of gathering information.

Police departments often work as agents of the courts, who attack by way of warrants or subpoenas. The subject of a warrant or subpoena may be ordered to keep silent about it.

Attacks by legislators may include declarations that keys must be escrowed, passing "Know Thy Customer" laws and identity card laws, and other measures usually taken with the public's interest in mind, but from an authoritarian point of view.

3.5 Hackers

Hackers will generally use search engines, Trojan horse software, and network monitoring (much like a sysadmin) to gather information about someone. Depending on their level of interest, they have also broken into credit reporting agencies, police computers, and other places with poor security to gather information.

3.6 National Intelligence

National Intelligence Agencies may operate wide net 'vacuum cleaner' operations designed to gather huge amounts of electronic information based on keywords, and who talks to whom. The Echelon system is reputed to do this. They may also engage in more targeted methods where they gather information from colleagues and acquaintances of people, or in technical attacks, where they use techniques such as Van Eck monitoring or hidden microphones to gather information.

3.7 Litigious Groups

There are a variety of organizations who, feeling their interests threatened, spend huge amounts of money threatening and filing lawsuits. This capability can allow them to determine email addresses in reply blocks. These lawsuits may need to be filed in a number of countries.

3.8 Organized Crime

Criminal organizations may attempt to either subvert the network or the privacy of a nym. This type of attacker is more likely to use physical violence for employee subversion, theft, or breaking and entering. On the other hand, maintain no illusions that organized criminals are unsophisticated thugs. In many cases, organized gangs are better funded and better equipped than police forces.

4 Attacks Against the Freedom System

4.1 A Few Eye Openers

We've said a lot of things over the last year and a half or so. Sometimes, we've described features that would be in Freedom that aren't there yet, or aren't going to be there. We want to let you know about some things that we don't think are problems, but might be surprising to those who expect Freedom to be an anonymity system, rather than a pseudonymity system.

1. HTTP referer and browser fields are left in place. We do this to allow Freedom to work with those web sites which break when we turn them off. This is much less surprising when you think of Freedom as a pseudonymity product, rather than an anonymity product. This creates a problem when you change nyms; if the referer points to a unique URL, then the site you're looking at can correlate that both your nyms are owned by the same user. (Also see 4.2.10)

2. A multi-part/mime signature from your mailer or browser can compromise a nym by signing a message with your 'real' identity and that of a nym. The Keyword Alert feature likely won't catch this because of the nature of X.509/PKCS#7 signature encoding.

3. It is very difficult to find information that has been arbitrarily encoded in outgoing data (e.g. information in compressed files, various file formats, etc.). Thus, the Keyword Alert only scans normal text in outbound flows. This is a compromise between the reality of a multi-format environment and our promise to deliver the Keyword Alert feature. Even if we tried to scan all possible formats we would inevitably fail. So, rather than trying very hard and giving you false confidence, we're realistic, and let you know what are the limits to this feature.

4. An attacker can see when you are using Freedom. The Freedom protocols allow you to assume a new identity when you browse, but someone who is watching the network links can see that you are logging into the Freedom Network by watching the packets. They can't tell what you're doing, but can see that you are logged in, and by counting packets and seeing how long you're online, may be able to make certain assumptions. (Counting and timing packets is possible today since traffic shaping and link padding do not offer strong security as implemented. See 4.3, 3 for more information.)

5. Mail users and Usenet news senders are not logged at the Freedom Mail Gateways (and Usenet gateway), however they could be logged in that the information is available to the mail gateway in the current mail protocol. We do currently log message sizes and times for to keep track of volume and aid debugging, but this information is stripped of email and pseudonym names.

6. Freedom does not offer anonymous NNTP news reading, mostly for business reasons – a full USENET news feed is expensive to maintain. The current recommended work-around is to anonymously browse a web based news source

such as dejanews.com. However, some users may use their ISP's news feed. Users who do this open themselves up to a correlation attack, as their ISP or an eavesdropper could record or even mark the messages they are reading, and then observe any pseudonymous replies they make and correlate the marked messages. In a future version of freedom we may re-sell access to a third-party NNTP server.

7. If you buy Freedom with a credit card, we store various data about you. It cannot be correlated to your nyms. Our privacy statement on this subject has been audited by TrustE, and is at http://www.zeroknowledge.com/alternate/policy.asp#store

8. If you forward mail sent to willshakespeare@freedom.net while logged in as francisbacon@freedom.net, you create an association that is hard to remove. The same issue appears with sending mail as romeo@freedom.net, mentioning things that only a Montague could know, or in other ways making it clear that you have knowledge that only a different persona has.

4.2 Active Attacks

1. Packages such as Back Orifice, WhoWhatWhere, NetBus, Systems Management Server, PCAnywhere, and other remote management tools totally compromise your privacy if the administrator so chooses. Freedom does not contain defenses against these, because they are inherent to Microsoft Windows, and we can not protect you against them. Anyone who can send you an attachment which you execute, or who can spoof one of your friends so that you trust an application sent in email, can execute this attack. We suggest keeping your anti-virus software up to date, and not running programs sent to you by email.

2. ActiveX, Javascript, VBScript, Java, and other executable content can allow an attacker to find information about you. There have been problems demonstrated with all of these systems. We expect that there will be more problems. We do not believe that it is possible to effectively filter them, and suggest that you turn them off. Anyone who runs a web site can exploit this problem.

3. Netscape's "What's Related" feature sends Netscape a complete history of your browsing, across all nyms and in non-private mode. We recommend you turn it off. Only Netscape, or people monitoring network traffic to Netscape, can exploit this problem.

4. If your mail tool is HTML enabled, and someone sends your nym a message containing an img= link, and you read that message without a nym selected, and if you allow the connection out, the attacker can correlate nym to IP address. Anyone with a web site can exploit this problem by sending you email. We suggest putting your nym email into separate folders, and only reading those folders while you are off-line, or using Freedom.

5. Nym key lookup responses are not signed. The data in the public key database which is returned is signed, but the response is not. This leads to a situation

where 'Nym not found' and 'Incorrectly formated request' messages can be forged. There are two scenarios in which nym lookups happen.

- Nyms lookup other nyms keys to send encrypted emails or to verify signatures on received emails. These are not vulnerable to the forged negative acknowledgement attack as the request is made over an authenticated encrypted route to the key query server.
- During the route create process the exit AIP sends the key query request to the key query server, but does not encrypt it. Therefore this lookup is vulnerable to fogered negative response attack. This would be a targetted denial of service attack.

  Exploiting this problem requires the ability to forge arbitrary packets on the Internet, to perform traffic analysis to figure out which packet you want to replay.

6. Link authentication is done poorly. We are not releasing details of how to implement this attack, but simply state that it is possible. This allows an attacker to insert packets, but to get those packets up to the user, they need to be able to understand how the authentication works at both the telescope and link layers. In general, the data inserted must be arbitrary; to insert chosen specific data is computationally infeasible (it would require the ability to cryptanalyze either 128 bit Blowfish or DH in real time.)

7. There is no link layer serialization, which allows packets to be replayed. To exploit this problem requires that you be able to insert packets into the client's link to the first hop, and read packets from arbitrary places on the Internet near other Freedom Servers. Note that if you can make a guess as to where the client might be surfing you only have to watch that one spot (and the client). One example of a common home page would be with Netscape users the default homepage is a fixed Netscape site, other examples would be popular homepage sites.

8. The current DH exchange lacks a nonce, has a race condition where the sides may misunderstand which key bytes are for whom, and there is extra data in the DH exchange that is sent in the clear (port numbers, time to live). This should result in nothing more than a DOS attack executable by someone who can forge packets.

9. If you have configured your DNS settings to search domains, the domains which you search will be exposed to the wormhole and its upstream DNS servers. Only someone running a Freedom network node or a DNS server that is searched by that Freedom node can exploit this to discover that someone searching a given domain is using a certain exit node. The DNS queries themselves (source IP address, etc) are anonymized.

10. If you are actively browsing the web when you change from one nym to another, then a web site (or someone monitoring the Internet) can see the HTTP referer field as a link from one nym to the next. Using a blank home page as your home

page, and returning to your home page before switching identities can minimize this.

11. Time synchronization is done through Zero-Knowledge. The stratum 1 time server for the Freedom Network is run by Zero-Knowledge, rather than encouraging stratum 1 and 2 servers all around the network. This is a design flaw we haven't corrected yet, because of the usefulness of synchronized time, and the effort to ensure our partners are using good time sources.

12. An attacker who can capture or emulate the "No such key" or "Malformed Request" responses from the Key Query Server or Nym server can splice them into TCP connections. Some of these connections are made over encrypted authenticated connections via the AIP network. However, connections are made in the clear if, when starting the client, the NIQS tells the client that there is a new AIP.

4.3 Passive Attacks

Not all of the following attacks are fully passive, but they all involve large amounts of backend processing that we expect only law enforcement and intelligence agencies would have the resources to engage in.

1. The signature keys for the system do not have planned periodic rekeying. The link keys are generated anew from a Diffie-Hellman exchange hourly (actually a mutually authenticated DH – see the Freedom System 2.0 Architecture white paper for details). The telescope keys are generated anew each time you create a route. But the signature keys are not periodically rekeyed, and that may open us to attack. In addition, the design calls for the link keys to be directionally different, and that is not currently done. (The same key encrypts data sent from A to B, and B to A.) Simply breaking into the server to steal the signature key will allow you to impersonate the Freedom Server by engaging in IP spoofing and sending fake signed requests. Doing this is roughly equivalent to continuing to exploit the compromised Freedom server, but is much more noticeable.

2. Key expirations are not checked. This is due to a bug in an earlier client which accidentally issued keys with very short expiration dates. Effectively the expiration field has been deprecated due to this bug. When we introduce expirations in the next protocol revision, the format version will be changed.

3. In the current version of the protocol there is no link padding, cover traffic or traffic shaping. It might be argued that one at minimum needs some of these counter-measures to defend against traffic analysis, but our initial analysis suggests that these counter-measures are *probably necessary, but certainly not sufficient*. This is because even if one does implement a combination of these counter-measures there remain a number of attacks, not significantly harder than attacking a system without these counter-measures. The main example is the packet round-trip timing related attacks, where the attacker passively observes or actively (and plausibly deniably) induces latency variations to uniquely identify the source of

a route. These remaining attacks are expensive in bandwidth utilization to defend against, and the counter measures greatly hinder performance. Consider that to defend against timing attacks, even as a first step one would need to start by padding round-trip times to get cover, reducing all round-trip times to worst case round-trip.

It is not clear how to defend against these attacks. We are researching how to exploit this class of attack to better understand the issues, some of which are quite closely tied to network semantics, and underlying TCP/IP vulnerabilities which are pragmatically unavoidable network properties. We hope this will help develop efficient counter-measures.

Wei Dai published an attack on the Freedom system, as described in the April 1999 white paper, "The Freedom Network Architecture." This attack was based on the nature of the traffic shaping system implemented in the prototype system at the time. Since traffic shaping is not enabled in the current version of the protocol, attacks on the previously planned mechanisms are not currently applicable. Dai's attack will be considered as part of our research into traffic analysis, and subsequent protocol revisions.

4. In Freedom 1.0 all data packets in the network were fixed sized. This lead to reduced useful bandwidth being available to the user. This is because the data packets are smaller than normal TCP streaming packets (which are typically 1500 bytes), in addition the acknowledgement packets in TCP streams are larger than usual because of the space overhead of padding them up to the fixed packet size, and there are more of them than usual as the data packets are smaller. The argument for fixed size packets is to hide information from a passive attacker.

   In Freedom there is no link padding, and the network is not synchronous. These two facts taken together mean that fixed size packets are not sufficient to prevent correlation attacks. Other attacks based on observing the effects of network congestion on links and observing timing correlations between client response time and observed data leaving the exit node allow a similarly powerful attacker to make correlations with or without fixed size packets. (Note an attacker can also create congestion with plausible deniability – just by using the network heavily over the target links.) Therefore fixed sized packets offer very limited additional traffic analysis resistance, but cost a lot in bandwidth and throughput. In freedom 2.0, starting with client version 2.1 we will enable variable sized packets.

5. In Freedom 1.0 the default number of hops for web and internet traffic was 3 hops. In Freedom 2.1 the default number of hops for web and internet traffic has been reduced to 1 hop. Users can adjust the number of hops for web browsing and internet use upwards with settings of 1, 2 or 3 hops to select their preferred security vs perfomance trade-off.

6. In Freedom 2.1 the number of hops for connections to the core servers has been reduced to 2 hops. The default number of hops for connections to the core servers was 3 hops in both Freedom 1.0 and 2.0. (The core servers include the freedom mail system server and the freedom servers providing key lookup and routing

information). Although for core access 3 hops offers more security than 2 hops, the difference in security is relatively small compared to the difference between 1 and 2 hops. There are a number of attacks based on correlating information entering and leaving the entry and exit nodes, and for most of these attacks the middle hop does nothing to resist the attack, other than widening the radius of reachable nodes.

Also the security of a connection is difficult to measure concretely, because the security offered depends on the attackers a user is concerned about, and the capabilities of those attackers. Beliefs about these parameters vary and are not easily measurable. Therefore it is difficult to know which attacks to place most weight on defending against.

A remaining question is whether 3 hop internet access used with supporting 2 hop core access offers balanced security, there are arguments for and against; ultimately it is a a trade-off between performance, reliability on the one hand and security on the other.

Connections to the core are not directly comparable to general internet connections because:

- core connections are end to end encrypted (there is no in the clear traffic between exit node and web server)
- there is better cover traffic for core connections because all users are accessing the same core, and they are accessing the core for the same reasons – the connections are encrypted and the purposes are either information lookup, or mail delivery related.
- all core nodes are operated by ZKS

For these reasons we chose 2 as the fixed number of hops for core access and a user selectable number of hops – between 1 and 3 hops – for web and other internet traffic as a good trade-off between performance, reliability and security.

See [3] for a more detailed analysis of the attacks on anonymity providing networks.

7. There is a family of attacks where the attacker takes data from the first hop and the last hop and then engages in various attacks against the middle one to find out more information. Data can be gathered about the the first hop by watching who connects to it. Data can be gathered from the last hop by watching many wormholes; this is more invasive the longer it goes on. Any hop can also be compromised by breaking into the system via an OS flaw, mis-configuration, etc. Note that some of these are mis-characterized as being passive so that the first-last attack family is all in one place.

   (a) The first variants of this attack are where the first node is noted by seeing a route create packet. The route may then be compromised by someone who can see the whole network and follow the route create, or the nym may be compromised by seeing which AIP does a nym lookup. The first variant of

this is the route-create traffic analysis attack, the second is the nym-lookup variant of the first-last attack. (The nym-lookup variant is enabled if Zero-Knowledge logs the nym lookups, or if the nym server is compromised, or by watching which AIPs send packets to the key query server.)

(b) The warrant variant of the first-last attack would, if Zero-Knowledge maintained logs, be to present a search warrant for the logs at a certain time. Zero-Knowledge does not maintain nym lookup logs, and has no capability of doing so.

(c) The sniping variant of the first-last attack is to replace statistical analysis with denial of service attacks on the links between AIPs, or the AIPs themselves. This requires an opponent who can cut Internet links 'at will,' and is willing to do so. We don't believe there are adversaries who can shut down network links at will, and are willing to reveal that capability, but some national intelligence agencies might be willing and able to do so.

(d) The stop-the-Internet variant involves shutting things down on a larger scale to see if the connection of interest survives. Again, we don't believe there are adversaries who can shut down the Internet at will, and are willing to reveal that capability, but some national intelligence agencies might be willing and able to do so.

(e) The traffic-mangling variant is when a malicious entry AIP damages packets from a targeted IP address. A collaborating exit node can note that it is receiving corrupt packets, and share that information out of band with an entry node (or group of malicious entry nodes) for time/source IP correlation.

8. The security lever can be slid upwards. There are good arguments that the opponents who can attack a 3 hop nym based on traffic analysis can do so more easily if you've ever used the nym over fewer hops, and thus, moving the lever towards "Optimize for security" is misleading, since you can't increase the security of a compromised nym.

4.4 Network Failure Attacks

1. If you connect to a web site that automatically refreshes itself (for example, http://www.cnn.com/), and your route through the Freedom network becomes unavailable, and the web site in question uses generated or otherwise unique URLs, there could be a correlation created to your real IP address by the web site. Freedom should block all network connections until you dismiss the error message Freedom displays. The best course of action, should this occur, is to close the browser, dismiss the dialog, create a new route, then re-launch the browser.

4.5 Archived Data Attacks

1. Forward-secrecy is used on all secured network connections, so there is limited value for an attacker using a warrant to obtain keys, as the communication keys

are re-keyed within a minimum of half an hour of use.

2. Freedom Activation Codes are tied to credit cards. We don't tie the Activation Codes to Nym Tokens. We have spent a lot of time to ensure these systems are separate from all our other corporate systems, as documented in "Untraceable Nym Creation on the Freedom 2.0 Network."

   The next protocol revision will use ZKS ecash technology based on Brands [1] to issue blind tokens which are redeemed for service. There still remains a timing correlation attack if payment is made via an identifying payment mechanism (such as a credit card) as some types of service are pseudonymous, and therefore there is an observable creation time. There are two ways we can address this problem: a free trial period where a user is able to pay at random time in the trial period, or by advising users to create pseudonyms after a delay to get cover from other nym creations. (Note that the latter advise only offers cover from other users who follow the advise; if few follow it limited cover is offered).

5 Notes for Analysts

1. Randomness: On Linux, we use /dev/urandom. On Windows we use Yarrow-160 [2]. In addition, we take random data from packet timings and content. For example on servers we mix in user chosen random IVs which are encrypted to the servers public key. In this way entropy is transfered from a typically high entropy environment client operating system to a low entropy environment – rack-mount server equipment with no input devices.

2. There is no passphrase strength bar. We use passphrase stretching, suggested minimum lengths, and salting. Salting is included to protect people against dictionary pre-computation attacks, which can then more cheaply attack multiple passphrases, such as might be 'harvested' from freedom.dat files from many users.

3. Linux lacks auditing capabilities. We can't audit file accesses, process creation, socket creation, and other activities which would allow us to better monitor security of AIPs and other nodes.

4. There is no secure memory on client or server. This relates closely to the auditing point above, and the sysadmin software point above (3.2). Secure solutions to this require some setuid memory access code, which has more problems than not securing memory, in our view.

5. Entry and exit AIPs can see where they are in the chain. Entry nodes can see their location by the half-authenticated link encryption, and exit nodes can see a nym signature block.

6. Signature verification is Freedom dependent. We have not published data structures to allow people to verify signatures independently. In addition, there are

places where signatures are removed from messages before they go out of the network, which is unfortunate, and enables forgeries.

7. A breakthrough in the analysis of either the discrete log problem, our pseudo-random number generators, or the bulk ciphers we use, would have dramatic impacts on the security of many fielded cryptographic systems, including Freedom. Similarly, construction of a large quantum computer would put a large dent in modern cryptographic practice.

8. Microsoft Windows isn't a secure Operating System. The vulnerability of Windows is a fundamental problem, and there are a whole variety of attacks, of which those like Back Orifice and ActiveX only scratch the surface.

9. IP options are not currently removed from packets. This could allow an attacker to distinguish the operating system the user is using based on it's IP options and semantics "signature".

10. The protocol design requires certain important keys be stored online. In conjunction with the lack of planned re-keying (4.3, 1), this has the potential to be a substantial problem.

6 Competitive Analysis

1. Mixmaster email offers outbound email privacy that is superior to that offered by Freedom, because there is no reply feature. The one-way nature of the system means there is no point of attack. There are several ways Mixmaster users can receive replies, including setting up reply-blocks, which offer security equivalent to Freedom, but is harder to use; and announcing that the nym reads a certain mailing list or newsgroup, which offers better security, but doesn't scale well to many nyms.

2. We are not aware of a web browsing system that offers a level of privacy and security equivalent even to Freedom with one hop, since Freedom offers a choice of operators, while the competing systems only offer a single operator who may be logging. We consider the tools that offer to protect you from Java, ActiveX, etc, to be unreliable, and believe that you should turn these things off for your protection.

3. We are not aware of a chat system that offers a level of privacy equal to using Freedom with one hop.

4. We are not aware of a telnet or ssh privacy solution that offers privacy comparable to Freedom with one hop.

5. We are not aware of a news posting solution that offers the ability to carry on a conversation with ease and privacy comparable to Freedom. Other solutions offer the ability to post anonymously, but not pseudonymously. With a persistent pseudonym you maintain an identity, and can send and receive email as that pseudonymous identity.

6. There are systems which require no installed code to run.

7 Plans for Improvement

This is a "planned actions" section. We are not committing to dates, versions, or even implementing these fixes; however, these are high level views of our current intent.

We plan to add planned re-keying in the very near term. We may move from key lookups to short lived certificates, or certificates pushed by the client together with certificate revocation lists, or a distributed database at or around the same time, to address the lookup attacks.

We intend to fix the protocol issues in the near term also.

We are researching traffic analysis and how to exploit the class of attacks based on timing and influencing network events which class of attack to better understand the issues. We hope this will help develop efficient counter-measures.

A Change History

November 30, 1999 Made the following changes:

4.1.9 added POP not private

4.2.9 DNS queries are anonymized, not pseudonymized.

4.2.10 motd is signed

4.2.13 added – internal

4.3.3 rudimentary traffic shaping survived to ship.

4.3.4.e traffic mangling

5.1 /dev/Urandom

5.7 Fixed some chain issues; these are fixed in the code, and have been removed from the paper.

6.1 updated Mixmaster; reply blocks are NOT a part of the sw

A Fixed Y2K bug in Change History section

November 23, '99 Released Initial Version.

References

[1] Stefan Brands, "Rethinking Public Key Infrastructure and Digital Certificates – Building in Privacy", PhD Thesis, Eindhoven University of Technology, 1999.

[2] J. Kelsey and B. Schneier and N. Ferguson, "Notes on the Design and Analysis of the Yarrow Cryptographic Pseudorandom Number Generator", Sixth Annual Workshop on Selected Areas in Cryptography, Springer-Verlag, 1999.

[3] Adam Back, Ulf Möller and Anton Stiglic, “Traffic Analysis Attacks and Trade-Offs in Anonymity Providing Systems”, To appear in the Information Hiding Workshop Proceedings, 2001.