Freedom 1.0 Security Issues and Analysis

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

November 23, 1999

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, but arise when the system is not properly used. This disclosure includes all known design or implementation flaws, as well places where various trade-offs made 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 of 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 most of 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 in 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 has been worrying them. All programmers also had a chance to review all of our white papers, and the most experienced have been forced to walk through 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 code reviews and walkthroughs, 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 such 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 achieved this. 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 intelligence agency to effectively, reliably, and on an ongoing basis, break the privacy which we offer with regard to web browsing, chat, or other interactive services.

If you use email reply blocks, then a series of warrants, delivered in reverse order, may lead law enforcement or lawyers to your real email address. Email reply blocks are created and turned on by default. Zero-Knowledge considers this acceptable for the time being because it is better than available alternatives, all of which except Cypherpunk and Mixmaster remailers[1] can be compelled to compromise your privacy with a single warrant. This is a targeted attack, as opposed to what an intelligence agency can achieve.

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.

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 tickled by the immature.

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

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.

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 to 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

Generally, police departments go after encryption keys to force data recovery, including identity information. This is usually done overtly, in the form of warrants or intimidation. The police also operate covertly through actions such as emissions monitoring and "dumpster diving." One cannot assume that all police actions are authorized or even legal. 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 be silent about it.

Attacks by legislatures include declaring privacy illegal, declaring that keys must be escrowed, passing "Know Thy Customer" laws, identity card laws, and other Orwellian measures that states often impose out of fear and ignorance.

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 crappy security to gather information.

3.6 National Intelligence

National Intelligence Agencies may operate large 'vacuum cleaner' operations designed to gather huge amounts of information based on keywords, and who talks to whom. The Echelon system is reputed to do this. They may also engage in more targeted attacks where they gather information from people around you, or 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 all pug-faced thugs. In many cases, organized gangs are better funded and 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 said features 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 is 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 unqiue URL, then the site you're looking at can correlate that your two nyms are owned by the same user. (Also see 4.2.11)

2. We break the POP AUTH command. The POP3 protocol includes both a PASS (password) command, and an AUTH (authenticate) command. The AUTH command is a stronger authentication method, which almost no one uses, and we break it. We found this late in the beta testing stage, and are not fixing it because of the complexity of the options negotiations it can lead into, and because it seems so few people are using it.

3. 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 text scanning box likely won't catch this because of the nature of X.509/PKCS#7 signature encoding.

4. 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 text scanner only scans normal text in outbound flows. This is a compromise between the reality of a multi-format environment and our promise to deliver text scanning. 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.

5. 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 are not turned on. See 4.3.3 for more information.)

6. Mail and Usenet news are logged at the Freedom Mail Gateways (and Usenet gateway). This is a *pseudononymous* system. We track the mail for debugging and spam control purposes. This does not lead to any way of correlating a nym to you.

7. If you buy Freedom with a credit card, we store various data about you. It can not be correlated to your nyms. Our privacy statement on this 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. ("Oh what tangled webs we weave, when first we practice to deceive" (If only we had tech-writers like that, you'd have more fun reading this document.))

4.2 Active Attacks

1. 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 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 owns 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 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 nym 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. This attack is mitigated by the fact that the request is encrypted and authenticated as it goes in and out. 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, and to do so within a one hour period so that the link keys are all correct.

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 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 to (Netscape's homepage?), you only have to watch that one spot (and the client).

8. The 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 pseudononymized.

10. Freedom Message of the Day feature does not check for a signature, or otherwise authenticate the message. The message of the day is passed in the clear over a TCP connection and is vulnerable to spoofing. This is a design flaw.

11. 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 popular web pages, such as http://www.freedom.net as your home page can minimize this.

12. 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.

4.3 Passive Attacks

Not all the following attacks are fully passive, but involve large amounts of backend processing that we expect only police and intelligence agencies could engage in.

1. The signature keys for the system are not rotated as planned. The code hasn't been fully QA'd and regression tested, and so the signature keys for the system are static. The link keys are generated anew from a Diffie-Hellman exchange hourly (actually a mutually authenticated DH – see the Freedom Network 1.0 Architecture white paper for details). The telescope keys are generated anew

each time you create a route. But the signature keys are not rotated, 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 related to the point above. There is a plan to expire and rotate keys, but its not going to be done until the code has been QA'd more thoroughly, and as such, expired keys can be used in some places.

3. Cover traffic and traffic shaping are not enabled. The use of cover traffic and traffic shaping over the network to complicate the task of traffic analysis has been planned and discussed. It is currently disabled for a variety of reasons, including engineering difficulty and load on the servers. This makes a variety of traffic analysis attacks much easier. Someone who can tap most of the Freedom network and run statistical analysis across it, can exploit these issues. We expect that only national intelligence agencies can do this.[2]

4. There is a family of attacks where the attacker takes data from the first hop, 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 hops 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.

   (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 attack is enabled if Zero-Knowledge logs the nym lookups, or if the nym server is compromised, or possibly by watching which AIPs send packets to the nym server. AIPs maintain anonymous routes to the core network information servers, so the nym-lookup variant is harder to implement.)

   (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.

[2] 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 on, attacks on the previously planned mechanisms are not relevant. We will certainly take it into account when we build new traffic shaping algorithms.

(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.

5. The padding of packets may absorb too much randomness. The outflow of randomness from the pool is not well controlled, and we may empty the pool too quickly to pad packets. An opponent with substantial cryptanalytic capability who can exploit the low rate at which random bytes are added to the pool could take advantage of this.

6. 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 (e.g., 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 relaunch the browser.

4.5 Archived Data Attacks

1. Reply blocks can be opened by a series of court orders. The court would first need to serve a warrant (do we have a subpoena/warrant policy?) on Zero-Knowledge, to which we will reply with encrypted data. If they order us to decrypt it, we will comply, and provide them with more encrypted data, perhaps with a little plaintext that points them to the next server to send a warrant to. This process will repeat until they get a POP email address to serve a warrant on. Someone needs to be able to get either a search warrant or a subpoena for this. Zero-Knowledge has a policy of [what are our public statements on this?]

2. Freedom Serial numbers are tied to credit cards. We don't tie the serial numbers to tokens. We have been audited by (still confidential) for compliance with this,

and they (will) attest to our compliance. We have spent a lot of time to ensure these systems separate from all our other corporate systems, as documented in "Untraceable Nym Creation on the Freedom Network."

5 Notes for Analysts

1. Randomness: On Linux, we use /dev/random. On Windows we use the Yarrow library. In addition, we take random data from packet timings and content. If this was our main source of data, there would be attacks on it. But it is simply something we stir into our pools.

2. MAIP to MAIP connections are not done through the anonymous cloud. It has been asserted that each hop of a Freedom mail connection is through a full 3 hop anonymous tunnel. This is not the case. Each mail hop is anonymized by the MAIP, segmented into appropriate packet sizes, and sent to the next MAIP.

3. There is no passphrase strength bar. We use passphrase stretching, suggested minimum lengths, and salting, but we do not implement a strength bar. Salting is included to protect people against dictionary attacks which 'harvest' the freedom.dat files from many users.

4. Beta serial numbers were sent unencrypted; the sending of serial numbers to some people may enable a forward looking attack against nyms created with those serial numbers, if the user gets the same nym in the real network. Production serial numbers are sent encrypted in a process described in "Untracable Nym Creation in the Freedom Network."

5. 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.

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

7. AIPs and MAIPs can see where in the chain they are. Various aspects of the data structures allow an AIP or MAIP to see how many servers are in the chain and what is their position in the chain.

8. 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.

9. 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.

10. Windows isn't a secure OS. 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.

11. IP options are not currently removed from packets. As there are currently no clients that are not running on Windows, we don't need to remove the Windows-identifying IP parameters.

12. The implemented key hierarchy requires certain important keys be stored online. In conjunction with the lack of key rotation 4.3.1, this has the potential to be a substantial problem.

6 Competitive Analysis

1. Mixmaster email offers outbound email privacy that is slightly superior to that offered by Freedom, since the reply-block feature is not on by default.

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 commiting to dates, versions, or even implementing these fixes, however, these are high level views of our current intent.

We plan to add key rotation and a new key hierarchy in the very near term. We may move from key lookups to a certificate revocation list or distributed database at or around the same time, to address the lookup attacks.

We intend to fix the protocol issues after the new key hierarchy, management, and rotation code is in place. Possibly in parallel with this, we intend to offer better mail response tools than reply-blocks (these are hard because the mail server becomes a 'last' in the first/last attacks.) We may support both reply blocks and another technology for a while.

It is likely tat only after we've replaced reply-blocks and upgraded the protocols will we be adding traffic shaping and cover traffic. This is because these tools are only useful against large, powerful opponents, and until we fix the other issues, there is little point in fixing this.

8 Change History

November 23, '99

Released Initial Version.