Internet Engineering Task Force (IETF)
Internet-Draft
Intended Status: Standards Track
Expires: April 30, 2015

Phillip Hallam-Baker
Comodo Group Inc.

October 27, 2014

Privacy Protected Security Considerations
draft-hallambaker-prismproof-req-01

Abstract

PRISM is reputed to be a classified US government program that involves covert interception of a substantial proportion of global Internet traffic. This document describe the security concerns such a program raises for Internet users and security controls that may be employed to mitigate the risk of pervasive intercept capabilities regardless of source.

Status of This Memo

Copyright Notice

Table of Contents

1. Requirements

PRISM is reputed to be a classified US government program that involves covert interception of a substantial proportion of global Internet traffic. While the precise capabilities of PRISM are unknown the program is believed to involve traffic and meta-data analysis and that the intercepts are obtained with the assistance of intermediaries trusted by Internet end users. Such intermediaries may or may not include ISPs, backbone providers, hosted email providers or Certificate Authorities.

Government intercept capabilities pose a security risk to Internet users even when performed by a friendly government. While use of the intercept capability may be intended to be restricted to counter-terrorism and protecting national security, there is a long and abundant history of such capabilities being abused. Furthermore an agency that has been penetrated by an Internet privacy activist seeking to expose the existence of such programs may be fairly considered likely to be penetrated by hostile governments.

The term 'privacy protected' is used in this series of documents to describe a communications architecture that is designed to resist or prevent all forms of covert intercept capability. The concerns to be addressed are not restricted to the specific capabilities known or suspected of being supported by PRISM or the NSA or even the US government and its allies.

2. Attack Degree

Some forms of attack are much harder to protect against than others and providing protection against some forms of attack may make another form of attack easier.

The degrees of attack that are of concern depend on the security concerns of the parties communicating.

2.1. Content Disclosure

Content disclosure is disclosure of the message content. In the case of an email message disclosure of the subject line or any part of the message body.

The IETF has a long history of working on technologies to protect email message content from disclosure beginning with PEM and MOSS. At present the IETF has two email security standards that address confidentiality with incompatible message formats and different key management and distribution approaches.

S/MIME and PGP may both be considered broken in that they reveal the message subject line and content Meta-data such as the time. This problem is easily addressed but at the cost of sacrificing backwards compatibility.

2.2. Meta Data Analysis

Meta Data is information that is included in a communication protocol in addition to the content exchanged, This includes the sender and receiver of a message, the time, date and headers describing the path the message has taken in the Internet mail service. Meta-data analysis permits an attacker to uncover the social network of parties that are in frequent communication with each other.

Preventing disclosure of meta-data is possible through techniques such as dead drops and onion routing but such approaches impose a heavy efficiency penalty and it is generally considered preferable to limit the parties capable of performing meta-data analysis instead.

The IETF STARTTLS extension to email permits the use of TLS to encrypt SMTP traffic including meta-data. However use of STARTTLS has two major limitations. First SMTP is a store and forward protocol and STARTTLS only protects the messages hop-by-hop. Second there is currently no infrastructure for determining that an SMTP service offers STARTTLS support or to validate the credentials presented by the remote server. The DANE Working Group is currently working on a proposal to address the second limitation.

2.3. Traffic Analysis

Analysis of communication patterns may also leak information about which parties are communicating, especially in the case of synchronous protocols such as chat, voice and video.

Traffic analysis of store and forward protocols such as SMTP is more challenging, particularly when billions of messages an hour may pass between the major Webmail providers. But clues such as message length may permit attackers more leverage than is generally expected.

2.4. Denial of Service

Providing protection against denial of service is frequently at odds with other security objectives. In most situations it is preferable for a mail client to not send a message in circumstances where there is a risk of interception. Thus an attacker may be able to perform a Denial of Service attack by creating the appearance of an intercept risk.

Whether the potential compromise of confidentiality or service is preferable depends on the circumstances. If critical infrastructure such as electricity or water supply or the operation of a port depends on messages getting through, it may be preferable to accept a confidentiality compromise over a service compromise even though confidentiality is also a significant concern.

2.5. Protocol Exploit

Many protocols are vulnerable to attack at the application layer. For example the use of JavaScript injection in HTML and SQL injection attacks.

A recent trend in Internet chat services is to permit the participants in a group chat to share links to images and other content on other sites. Introducing a link into the chat session causes every connected client to retrieve the linked resource, thus allowing an attacker with access to the chat room to discover the IP address of all the connected parties.

3. Attacker Capabilities

Some forms of attack are available to any actor while others are restricted to actors with access to particular resources. Any party with access to the Internet can perform a Denial of Service attack while the ability to perform traffic analysis is limited to parties with a certain level of network access.

A major constraint on most interception efforts is the need to perform the attack covertly so as to not alert the parties to the fact their communications are not secure and discourage them from exchange of confidential information. Even governments that intentionally disclose the ability to perform intercepts for purposes of intimidation do not typically reveal intercept methods or the full extent of their capabilities.

3.1. Passive Observation

Many parties have the ability to perform passive observation of parts of the network. Only governments and large ISPs can feasibly observe a large fraction of the network but every network provider can monitor data and traffic on their own network and third parties can frequently obtain data from wireless networks, exploiting misconfiguration of firewalls, routers, etc.

A purely passive attack has the advantage to the attacker of being difficult to detect and impossible to eliminate the possibility that an intercept has taken place. Passive attacks are however limited in the information they can reveal and easily defeated with relatively simple cryptographic techniques.

3.2. Active Modification

Active attacks are more powerful but are more easily detected. Use of TLS without verification of the end-entity credentials presented by each side is sufficient to defeat a passive attack but is defeated by a man-in-the-middle attack substituting false credentials.

Active attacks may be used to defeat use of secure after first contact approaches but at the cost of requiring interception of every subsequent communication.

While many attackers have the ability to perform ad-hoc active attack only a few parties have the ability to perform active attack repeatedly and none can expect to do so with absolute reliability.

A major limitation on active attack is that an attacker can only perform an active attack if the target is known in advance or the target presents an opportunity that would compromise previous stored communications.

3.3. Cryptanalysis

Many parties have the ability to perform cryptanalysis but government cryptanalytic capabilities may be substantially greater.

3.4. Kleptography

Kleptography is persuading the party to be intercepted to use a form of cryptography that the attacker knows they can break. Real life examples of kleptography include the British government encouraging the continued use of Enigma type cryptography machines by British colonies after World War II and the requirement that early export versions of Netscape Navigator and Internet Explorer use 40 bit symmetric keys.

3.4.1. Covert Channels in RSA

One form of kleptography that is known to be feasible and is relevant to IETF protocols is employing an RSA modulus to provide a covert channel. In the normal RSA scheme we choose primes p and q and use them to calculate n = pq. But the scheme works just as well if we choose n' and p and look for a prime q in the vicinity of n'/p then use p and q to calculate the final value of n. Since q ~= n'/p it follows that n' ~= n. For a 2048 bit modulus, approximately 1000 bits are available for use as a covert channel.

Such a covert channel may be used to leak some or all of the private key or the seed used to generate it. The data may be encrypted to avoid detection.

3.4.2. Covert Channels in TLS, S/MIME, IPSEC

Similar approaches may be used in any application software that has knowledge of the actual private key. For example a TLS implementation might use packet framing to leak the key.

3.4.3. Covert Channels in Symmetric Ciphers

A hypothetical but unproven possibility is the construction of a symmetric cipher with a backdoor. Such an attack is far beyond the capabilities of the open field. A symmetric cipher with a perfect backdoor would constitute a new form of public key cryptography more powerful than any known to date. For purposes of kleptography however it would be sufficient for a backdoor to limit the key space that an attacker needed to search through brute force or have some other limitation that is considered essential for public key cryptography.

3.4.4. Covert Channels in ECC Curves

Another hypothetical but unproven possibility is the construction of a weak ECC Curve or a curve that incorporates a backdoor function. As with symmetric ciphers, this would require a substantial advance on the public state of the mathematical art.

3.4.5. Unusable Cryptography

A highly effective form of kleptography would be to make the cryptographic system so difficult to use that nobody would bother to do so.

3.5. Lawful Intercept

Lawful intercept is a form of coercion that is unique to government actors by definition. Defeating court ordered intercept by a domestic government is outside the scope of this document though defeating foreign lawful intercept requests may be.

While the US government is known to practice Lawful Intercept under court order and issue of National Security Letters of questionable constitutional validity, the scope of such programs as revealed in public documents and leaks from affected parties is considerably more restricted than that of the purported PRISM program.

While a Lawful Intercept demand may in theory be directed against any of the intermediaries listed in the following section on subversion or coercion, the requirement to obtain court sanction constrains the number and type of targets against which Lawful Intercept may be sought and the means by which it is implemented. A court is unlikely to sanction Lawful Intercept of opposition politicians for the political benefit of current office holders.

3.6. Subversion or Coercion of Intermediaries

Subversion or coercion of intermediaries is a capability that is almost entirely limited to state actors. A criminal organization may coerce an intermediary in the short term but has little prospect of succeeding in the long term.

3.6.1. Physical Plant

The Internet is at base a collection of data moving over wires, optical cables and radio links. Every form of interconnect that is a practical means of high bandwidth communication is vulnerable to interception at the physical layer. Attacks on physical interconnect require only a knowledge of where the signal cables are routed and a back hoe.

Even quantum techniques do not necessarily provide a guarantee of security. While such techniques may be theoretically unbreakable, the physical realization of such systems tend to fall short. As with the 'unbreakable' One Time Pad, the theoretical security tends to be exceptionally fragile.

Attacks on the physical plant may enable high bandwidth passive intercept capabilities and possibly even active capabilities.

3.6.2. Internet Service Providers

Internet Service Providers have access to the physical and network layer data and are capable of passive or active attacks. ISPs have established channels for handling Lawful Intercept requests and thus any employee involved in an intercept request that was outside the scope of those programs would be on notice that their activities are criminal.

3.6.3. Router

Compromise of a router is an active attack that provides both passive and active intercept capabilities. such compromise may be performed by compromise of the device firmware or of the routing information.

3.6.4. End Point

Compromise of Internet endpoints may be achieved through insertion of malware or coercion/suborning the platform provider.

3.6.5. Cryptographic Hardware Providers

Deployment of the 'kleptography' techniques described earlier requires that the attacker be capable of controlling the cryptographic equipment and software available to the end user. Compromise of the cryptographic hardware provided is one means by this might be achieved.

3.6.6. Certificate Authorities

Certificate Authorities provide public key credentials to validated key holders. While compromise of a Certificate Authority is certainly possible, this is an active attack and the credentials created leave permanent evidence of the attack.

3.6.7. Standards Organizations

Another route for deployment of cryptography would be to influence the standards for use of cryptography although this would only permit the use of kleptographic techniques that are not publicly known.

Another area of concern is that efforts to make strong cryptography usable through deployment of key discovery infrastructure or security policy infrastructure may have been intentionally delayed or discouraged. The chief security failure of the Internet today is that insecurity is the default and many attacks are able to circumvent strong cryptography through a downgrade attack.

4. Controls

Traditionally a cryptographic protocol is designed to resist direct attack with the assumption that protocols that provide protection against targeted intercept will also provide protection against pervasive intercept. Consideration of the specific constraints of pervasive covert intercept demonstrates that a protocol need not guarantee perfect protection against a targeted intercept to render pervasive intercept infeasible.

One of the more worrying aspects of the attempt to defend the legality of PRISM program is the assertion that passive intercept does not constitute a search requiring court oversight. This suggests that the NSA is passively monitoring all Internet traffic and that any statement that a citizen might make in 2013 could potentially be used in a criminal investigation that began in 2023.

At present Internet communications are typically sent in the clear unless there is a particular confidentiality concern in which case techniques that resist active attack are employed. A better approach would be to always use encryption that resists passive attack, recognizing that some applications also require resistance to active attacks.

4.1. Confidentiality

Encryption provides a confidentiality control when the symmetric encryption key is not known to or discoverable by the attacker. Use of strong public cryptography provides a control against passive attacks but not an active attack unless the communicating parties have a means of verifying the credentials purporting to identify the parties.

4.1.1. Perfect Forward Secrecy

One of the main limitations of simple public key exchange schemes is that compromise of an end entity decryption key results in compromise of all the messages encrypted using that key. Perfect Forward Secrecy is a misnomer for a technique that forces an attacker to compromise a separate private key for every key exchange. This is usually achieved by performing two layers of public key exchange using the credentials of the parties to negotiate a temporary key which is in turn used to derive the symmetric session key used for communications.

Perfect Forward Secrecy is a misnomer as the secrecy is not 'perfect', should the public key system used to identify the principals be broken, it is likely that the temporary public key will be vulnerable to cryptanalysis as well. The value of PFS is not that it is 'perfect' but that it dramatically increases the cost of an attack to an attacker.

4.2. Policy, Audit and Transparency

The most underdeveloped area of internet security to date is the lack of a security policy infrastructure and the audit and transparency capabilities to support it.

4.2.1. Policy

A security policy describes the security controls that a party performs or offers to perform. One of the main failings in the Internet architecture is that the parties have no infrastructure to inform them of the security policy of the party they are attempting to communicate with except for the case of Certificate Policy and Certificate Practices Statements which are not machine readable documents.

A machine readable policy stating that a party always offers a minimum level of security provides protection against downgrade attack.

4.2.2. Audit

Audit is verifying that a party is in compliance with its published security policy. Some security policies are self-auditing (e.g. advertising support for specific cryptographic protocols) others may be audited by automatic means and some may require human interpretation and evaluation.

4.2.3. Transparency

A security policy is transparent if it may be audited using only publicly available information.

An important application of transparency is by trusted intermediaries to deter attempted coercion or to demonstrate that a coercion attempt would be impractical.

Author's Address

Phillip Hallam-Baker
Comodo Group Inc.

philliph@comodo.com