Examining Proxies to Mitigate Pervasive Surveillance

Eliot Lear
Barbara Fraser

Abstract

The notion of pervasive surveillance assumes that it is possible for an attacker to have access to all links and devices between end points, as well as end points themselves. We examine this threat is some detail with an eye toward whether trusted intermediaries can provide relief from the attack. We go on to examine the costs associated with the various remediation methods. In at least one case, we challenge the notion that one should encrypt absolutely everything in all cases, as was implied in at least one threat analysis. Finally we summarize in a set of four principles that should be considered in future work.

1. Introduction

Recent disclosures have highlighted a new form of threat to Internet Security. Various attacks have been undertaken that seek to examine content and so-called meta-information between devices, perhaps linking them to specific individuals. In order to understand how best to protect against a threat one must understand what information is both available and of some value to the attacker. To begin this discussion we examine a typical IP packet. It contains, amongst other things, a source and destination address pair, a protocol ID, a time to live, source port, destination port, and some transport-layer information in the header. It also contains application data[1] [2] [3]. Each of these pieces of information provides some value to an attacker. An IP address may be correlated to a location, if not a specific individual or a service. Put another way, if an attacker has access to the path between two participants, without looking at the application content, she can with some degree of certainty determine who the participants in a communication are, where they are, when they have communicated (via timestamp), and what application and what service they are using. This assumes an end-to-end view of the world. Network Address Translation (NAT)[4] obscures the precise identity of a device, but may still provide a reasonable approximation of location. The attacker may also be in a position to determine location accuracy by the owner of an IP address. For instance, a block assigned to a mobile network provider can be assumed to have a fairly wide range – perhaps worldwide, whereas a block assigned to a fixed access network may be localized to a city or even a part of a city. Furthermore, the block itself may provide hints as to appropriate analysis to further localize the end point (such as analysis of exfiltrated 3G/4G information). Transport Layer Security (TLS)[5] on its own does not protect against these forms of attack. IPsec[6] in transport mode only ameliorates the problem by removing transport and application protocol information.

2. A Potential Role and Challenge for Proxies

An obvious approach to reducing the attacker's ability to observe all traffic is to aggregate it at certain trusted points. shim or L3 proxies such as SOCKS help obscure information on either side. That is, if Bob is talking through a proxy to Sally, an observer at one location may only be able to glean information that either Sally **or** Bob is communicating, but that they are not communicating with each other. If the

observer is present on both sides of the proxy, then it will glean that both Bob and Sally are communicating, but not *necessarily* to each other. Rather it is the proxy that is identified. However, there are several limitations to this protection:

1. If TLS is not employed on both sides, then an attacker observing on both sides will assuredly establish that Bob and Sally are communicating with each other, as well as the precise nature of their communication. The mitigation is to use TLS at both ends, and in the middle.
2. Even if TLS is employed, if only Bob and Sally are communicating with the proxy, then it can easily be inferred that they are communicating with each other. The mitigation is to aggregate sufficient clients such that the probability of only a small number of devices making use of the proxy is small.
3. Even if many systems are communicating with the proxy, examination of packet timing and size over time will reveal that Bob and Sally are communicating. Sending additional dummy content partially mitigates this attack.
4. Finally, the proxy itself may be compromised by an attacker, in which case Bob, Sally, and everyone else using the proxy are revealed to the attacker. There is no mitigation against this attack.

We'll discuss the costs of the approach as well as mitigations below.

In the case of SOCKS, each and every application must be modified to use the SOCKS proxy application gateway. This will be true of other non-transparent application layer proxies, of course. In fact many common applications defined within the IETF provide for proxies of some form within their architecture. The benefit of SOCKS or something similar is that all layers above the IP layer may be protected. SOCKS' drawback is that it has limited scale capability and is difficult to implement for different transports, especially when sessions do not exist. No consideration has ever been given to MP-TCP, for instance.

Working within layer three, a simple alternative would be to have an IPsec-based tunnel mode proxy service. This provides many of the benefits of SOCKS, but without the headaches of having to modify applications.

Moving further into more recent approaches, ToR networks work to obscure both source and destination addresses at different points in the network through the use of Onion Routing.[7] They too serve to address the precise threat discussed in Section 1. ToR-based networks offer the additional protection that they may route traffic amongst themselves for the purposes of limiting traffic analysis. However, if the attacker is observing both Bob and Sally, or observing the proxy, a timing analysis attack in a low latency environment will almost certainly be effective.[8] [9]

The benefit of using specific application aware proxies is that they can perform other activities based on the knowledge of the application. For instance, HTTP proxies may examine web traffic for malware, cache objects, while SMTP MTAs might reject connections from known spammers. The price paid for using specific application layer gateways is that the service being used may be revealed to the observer. Whether that information is valuable will depend on many factors, only some of which can be predicted.

It should be noted that large aggregated services such as Gmail act as a proxy in as much as they serve as a clearinghouse between senders and receivers of vast volumes of mail. Encrypted Email traffic is only subject to limited traffic analysis

because people often read email via protocols that are different from how mail is transmitted. Hence, size and timing analysis may prove more difficult. That is- given end-to-end encryption, an observer may be able to see that Bob sent a message to a large mail site, but it may not be able to detect either that Sally was the intended recipient or that even if it presumed Sally was, that she has read that specific message at any given time.

Email is store and forward, but there are other ways to proxy, especially at large scale: Cisco's Scansafe is a large proxy. And companies such as Amazon, Google, and Opera all provide proxy functionality for their web browsers which provide some security for their users, faster page rendering, and advertising and tracking value to those companies.

An HTTP2-based proxy may also serve as an effective blurring or blinding function if both Bob and Sally are communicating with others, and the proxy itself is heavily used.

3. Existing Usage and Economic Aspects of Proxies

In this section we will examine costs of running proxies, who may be perceived to benefit from that proxy, and whether incentives in fact would align to allow them to operate.

It should be noted that many proxies already exist. Both enterprises and service providers aggregate Email for purposes of central storage, spam and malware mitigation, and efficiency for mobile devices. In the case, of web proxies, the situation is considerably more complex. Enterprises often make use of web-based proxies for the purposes of malware prevention and policy enforcement. In addition, NAT-based firewalls offer at least some limited protection by aggregating large numbers of IP addresses into a relatively small pool. N.B., any protection based on the IP address can be muted through timing attacks, TTL examination, and for that matter, application content if the stream is not encrypted.

Service providers make use of different sorts of proxies. Some provide proxy in order optimize bandwidth consumption, while others may use proxies for other activities. Some use of Carrier Grade Network Address Translation (CGN) performs the same NAT function from which an enterprise benefits, but within the SP environment. It should be noted that in this case, the economic value of the NAT is not tied to privacy, but rather address space utilization. As such, deployment of IPv6 may reduce usage of CGN, without any regard to privacy. A standalone "privacy proxy" is a proxy that exists for the sole purpose of obscuring traffic and its content. This may be a rare breed, as we see below.

Let us presume, for sake of discussion that the service provider is itself not trusted by the end user. Can the end user go to some other "proxy service"? This leads to a number of issues:

1. Such proxy services will essentially consume bandwidth both incoming and outgoing that is roughly equivalent to the customer's use, absent caching functions. Someone will have to pay for that bandwidth. The customer might pay for this service directly, or through other means such as allowing ad insertion. This may be particularly problematic for ToR users, as service providers themselves may see no value in offering ToR for free. Use of onion

routing technology in other aspects may prove useful if someone is willing to foot the bill for the bandwidth consumed by proxies.

2. Related, as much as the world is mobile, use of proxies can introduce stretch into the network. That is, the path taken between two ultimate endpoints must go through a point that is not within the most efficient path of communication. Stretch increases bandwidth costs due to its transit effects, as well as the potential for increased delay. To mitigate this problem, it would be necessary to have some secure means to determine a proxy that is in fact near to the customer. Any middle box such as the proxy adds both latency and jitter, which reduce the user experience for real time communications such as interactive voice calls.
3. If uptake of such a service is small, its use may on its own bring the customer to the attention of the attacker.

There may be other challenges as well. A proxy will reside within some jurisdiction where a government may have jurisdiction and require disclosure of communications.

4. Principles

1. End-to-end encryption is necessary but not sufficient

In examining the above approaches, the reader is reminded that a substantial amount of information is revealed without ever having to look at the content of a packet. On the other hand, when a packet is **not** encrypted, much of the information that might be found in the header can **also** be found in its content, and more. Thus, encryption is a necessary, but not sufficient, measure to counter pervasive surveillance.

2. Parties must decide who outside the communication they will trust

In considering how so-called “meta-information” is obscured, participants attempting to use other parties to hide their path must to some extent trust those other parties. Those other parties are subject to attack or to the laws of a particular jurisdiction. As a result, parties must choose who they wish to trust. In most cases today, one party or the other employs a third party. For example, an HTTP or SIP client makes use of a proxy, or the server makes use of a reverse-proxy service, the model being that these proxies are acting as agents of one party or the other. However, both parties are, in effect, extending trust. Furthermore, lawful intercept regulations are not likely to abate. The same rules that apply today will apply in the future, and access service providers will be required to comply.

3. Laws of Economics must be observed

Any service that offers the blinding or blurring function such as ToR or a VPN proxy must be economically sound. Someone has to pay for the infrastructure it takes to establish any overlay network. If the service is built into an access service provider’s infrastructure, the access service provider must have some economic incentive to maintain it.

4. Blinding or blurring may work in limited scenarios

Finally, blinding and blurring are subject to timing and traffic analysis attacks. The size of transactions may provide limited protection through either meaningful or meaningless transformations. Email provides an example of a meaningful transformation. Random padding provides an example of meaningless transformation. A random amount of padding, however, is unlikely to hide

transactions involving large amounts of data, such as interactive voice or video, or receipt of large files if the observer is near one of the end points.

5. Conclusions

This paper has examined how encryption alone may not provide sufficient protection against pervasive monitoring, and how some form of proxy technology may provide assistance, but at a cost. We conclude that there will be challenges in developing standalone proxy services because of the cost in running them. If sufficient numbers of customers value privacy, then the cost may sustain a market. But if that market doesn't materialize in a substantial way, those who do make use of standalone "privacy proxies" make mark themselves as targets. However, if bundled to other services that provide benefit to either service providers or enterprises, a privacy proxy service may all serve the needs of interested consumers.

To accomplish such a service in the market place, timing attacks and secure proxy discovery that avoids stretch must be further explored.

6. Acknowledgments

The authors wish to thank Dan Wing, Cullen Jennings, and Eric Vyncke for their helpful comments.

1 Postel, J. Ed., “Internet Protocol”, RFC 791 (also STD 5), September, 1981.
2 Postel, J. Ed, “User Datagram Protocol”, RFC 768 (also STD 6), August, 1980.
3 Postel, J. Ed, “Transmission Control Protocol”, RFC 793 (also STD 7), September, 1981.
4 SriSuresh, P., Evegang, K., “Traditional IP Network Address Translator (Traditional NAT)”, RFC 3022, January, 2001.
5 Dierks, T., Rescrorla, E., “The Transport Layer Security (TLS) Protocol Version 1.2”, RFC 5246, August 2008.
6 Kent., S., “IP Encapsulating Security Payload (ESP)**”,** RFC 4303, December 2005.
7 http://en.wikipedia.org/wiki/Onion_routing
8 Shmatikov, Wang; Ming-Hsiu Vitaly (2006). "Timing analysis in low-latency mix networks: attacks and defenses". *Proceedings of the 11th European conference on Research in Computer Security*. ESORICS'06: 18–33
9 http://blog.ioactive.com/2012/02/ssl-traffic-analysis-on-google-maps.html