

Talking to Yourself for Fun and Profit

Lin-Shung Huang*, Eric Y. Chen*, Adam Barth[†], Eric Rescorla[‡] and Collin Jackson*

*Carnegie Mellon University, {linshung.huang, eric.chen, collin.jackson}@sv.cmu.edu

[†]Google, adam@adambarth.com

[‡]RTFM, ekr@rtfm.com

Abstract—Browsers limit how web sites can access the network. Historically, the web platform has limited web sites to HTTP, but HTTP is inefficient for a number of applications—including chat and multiplayer games—for which raw socket access is more appropriate. Java, Flash Player, and HTML5 provide socket APIs to web sites, but we discover, and experimentally verify, attacks that exploit the interaction between these APIs and transparent proxies. At a cost of less than \$1 per exploitation, our attacks poison the proxy’s cache, causing all clients of the proxy to receive malicious content supplied by the attacker. We then propose a modification of the HTML5 WebSocket protocol that resists these (and other) attacks. The WebSocket working group has adopted a variant of our proposal.

I. INTRODUCTION

Browsers restrict how web applications can interact with the network by enforcing a number of security invariants on their use of the user’s network connection. These restrictions are essential to the core security guarantee of the web security model: users can safely visit arbitrary web sites and execute scripts provided by those sites. Generally speaking, browsers permit web applications to send well-formed HTTP requests to arbitrary network locations (with a handful of important restrictions) but prevent them from reading back the response unless the server opts in via some mechanism.

A number of plug-ins relax these restrictions. For example, both Java and Flash Player provide a mechanism for web applications to open raw socket connections. Of course, unrestricted raw socket access to the network would be disastrous for security. An attacker could use such a facility to wreak havoc with any network service that relies on IP source addresses or network connectivity for security (e.g., network devices behind a firewall). Rather than allowing unrestricted socket access, both Java and Flash Player limit web applications to opening sockets to servers that have *consented* to such connections.

Java and Flash Player use different consent protocols. Java uses a trivial “consent” protocol whereby Java bytecode is implicitly authorized to open socket connections to the IP address from which it originated. Flash Player, by contrast, requires the server to supply a policy file over a specific port that explicitly authorizes socket connections to a set of port numbers. Although these protocols are widely deployed in browsers, the protocols themselves have seen only modest amounts of security analysis. Recently, these protocols were shown to be vulnerable to DNS rebinding attacks [1], whereby the consent was scoped to a host name rather than an IP address, letting the attacker transfer his or her consent to another network endpoint.

In our study, we show that the consent protocols used by browsers today are vulnerable to attack in certain network configurations involving network intermediaries, specifically transparent proxies. Unlike traditional HTTP proxies, which are explicitly configured and known to the client, transparent proxies insert themselves into the transport path (e.g., by acting as the network’s default gateway or as a bridge) and then act as proxies without the client’s knowledge. Such proxies are common in traffic filtering applications but also can serve as network accelerators or proxy caches. Although colloquially referred to as “transparent” proxies, these proxies are more accurately termed “intercepting” proxies because, as we show in this paper, they are not quite as transparent as their deployers might wish.

Unfortunately, these transparent proxies often forward the server’s consent without understanding its semantics. When a server provides a Flash policy file authorizing a SWF to connect to the server’s IP address on port 80, Flash Player will allow the SWF to open a raw socket connection to the server, not aware that the SWF is actually talking to a transparent proxy instead of the server itself. Once the attacker has opened a socket to the proxy server, the type of misdeeds the attacker can perform depend on details of how the proxy behaves.

Auger [2] describes how an attacker can leverage transparent proxies to establish connections with any host accessible by the proxy. We introduce new attacks that can poison the proxy’s cache for an arbitrary URL, causing *all* users of the proxy to receive the attacker’s malicious content instead of the honest server’s content. The conditions required for such an exploit are relatively precise (i.e., a specific class of proxy behavior). To determine whether these conditions actually arise in practice, we conduct experiments on the Internet to see what fraction of Internet users are vulnerable to these attacks by running an advertisement that mounted the attacks against servers in our laboratory. We found that 3,152 of 51,273 users (6.1%) in our study were vulnerable to Java-based IP hijacking attacks and 2,109 of 30,045 (7%) were vulnerable to Flash Player-based IP hijacking attacks. Furthermore, 53 of 30,045 (0.18%) users in our study were vulnerable to Java-based cache poisoning attacks and 108 of 51,273 (0.21%) were vulnerable to Flash Player-based cache poisoning. We believe that such attacks are critical, since every successful cache poisoning attack would also affect all users of the vulnerable proxy (potentially the entire enterprise), causing further impact beyond our raw measurements. Our experiments demonstrated one successful cache poisoning attack per \$0.93 spent on advertisements.

Raw socket APIs let web applications provide functionality that is difficult to provide with only HTTP networking APIs. Rather than simply recommending that raw socket access be removed from the web platform, we study the question of how to design a consent protocol that is robust to oblivious intermediaries. As a starting point, we consider the protocol for HTML5’s socket API, WebSockets [3], [4]. WebSockets uses an “in-band” consent protocol whereby the browser exchanges messages with the server over a socket before handing the socket over to the web application.

We show, empirically, that the current version of the WebSocket consent mechanism is vulnerable to proxy cache poisoning attacks. Even though the WebSocket handshake is based on HTTP, which should be understood by most network intermediaries, the handshake uses the esoteric “Upgrade” mechanism of HTTP [5]. In our experiment, we find that many proxies do not implement the Upgrade mechanism properly, which causes the handshake to succeed even though subsequent traffic over the socket will be misinterpreted by the proxy.

Building upon our analysis and empirical measurements on strawman protocols, we propose improving the WebSocket protocol by randomizing the attacker-controlled bytes sent on the wire. By encrypting the bytes sent on the wire using a stream cipher with a fresh random nonce for each protocol frame, the attacker cannot choose arbitrary bytes on the wire, making it difficult to confuse the receiver into performing undesirable actions. Our essential insight is that protocol designers should consider how attackers can manipulate these protocols to exploit network intermediaries that unintentionally proxy the consent of the remote server without understanding its semantics.

Contributions. The main contributions of this paper can be summarized as follows:

- We introduce a new class of attacks that poisons the HTTP caches of transparent proxies via socket APIs in Flash Player and Java, causing malicious content of the attacker’s choice to be served by the proxy to all of its users. Our experiments verify that roughly 7% of Internet users are vulnerable to Auger’s IP hijacking attacks, while 0.2% are vulnerable to our cache poisoning attacks.
- We demonstrate these attacks on HTML5 WebSocket strawman protocols. We propose improving the WebSocket protocol by encrypting the bytes sent on the wire using a stream cipher, making the payload data appear random to network entities that are oblivious to WebSockets. In response to our suggestion, the WebSocket protocol working group adopted a variant of our proposal that masks attacker-controlled bytes with XOR “encryption” instead of a stream cipher.

Organization. The rest of this paper is organized as follows. Section II explains the existing network access mechanisms in browsers. Section III details our attacks on Flash Player and Java, including our experimental verification. Section IV demonstrates attacks on HTML5 WebSocket strawman protocols and presents our proposal. Section V places our work in the context of related work. Section VI concludes.

II. BACKGROUND: NETWORK ACCESS IN THE BROWSER

In this section, we review the network access mechanisms browsers provide to web applications in the context of a threat environment. Consider a network topology in which the user connects to the Internet via a transparent proxy, as is common in enterprise networks. The transparent proxy intercepts outbound HTTP requests, perhaps to monitor employee network access, to enforce a security policy, or to accelerate web traffic.

In this scenario, we wish the browser to enforce a set of security policies that prevent malicious web sites from interacting arbitrarily with other hosts from the client’s IP address. Our assumption is that the user visits the malicious web site, that the browser properly enforces its security policy, and that the attacker has no direct control over the network intermediaries. The relevant question, then, is what security policy should the browser enforce on the malicious web site’s network access?

A. Same-Origin Policy

One natural response to the threat of web attackers is to simply forbid web applications running in the browser from communicating with any server other than the one hosting the application. This model, called the *same-origin policy*, was first introduced for Java applets. Java was originally designed as a general purpose programming language and so, unsurprisingly, offers generic networking primitives, including an API that lets the programmer request the virtual machine to open a raw socket to an arbitrary network address and port. If the virtual machine fulfilled these requests unconditionally, these API would be extremely dangerous. For this reason, Java allows network connections only to the source of the Java bytecode.¹ The policy appears, *a priori*, safe; how much harm can you cause if you’re talking only to yourself?

Unfortunately, Java’s notion of “source” has proved to be quite problematic. One natural definition of “source” is to simply compare host names, but there is no guarantee that the same host name will always be bound to servers controlled by the same entity. In particular, if the Java virtual machine does its own name resolution, then the system becomes vulnerable to DNS rebinding attacks [1], [6]. In these attacks, the victim visits the attacker’s web site (e.g., *attacker.com*) while the attacker’s DNS server responds to user’s initial DNS query with an A record pointing to the attacker’s server but with a short time-to-live. The client downloads the Java applet, which then opens a socket to *attacker.com*. Because the DNS response has expired, the Java virtual machine resolves the host name again, but this time the attacker serves an A record pointing to the target server, letting the applet (which is under the attacker’s control) open a socket to the target server from the client’s IP address. DNS rebinding attacks have been known for a long time and are addressed by basing access control decisions on the IP address rather than the host name, either directly by checking against the IP address (as in Java) or by

¹These restrictions do not apply to signed applets which the user has accepted. Those applets have the user’s full privileges.

pinning, forcing a constant mapping between DNS name and IP address regardless of the time-to-live of the DNS response.

B. Verified-Origin Policy

Unfortunately, the same-origin policy, strictly construed, is quite limiting: many web application developers wish to communicate with other web sites, for example to incorporate additional functionality or content (including advertisements). Allowing such communication is unsafe in the general case, but the browser can safely allow communication as long as it verifies that the target site consents to the communication traffic. There are a number of Web technologies that implement this *verified-origin policy* [7]:

1) *Flash Cross-Domain Policies*: Prior to letting a SWF open a socket connections to a server, Flash Player first connects to the site and fetches a cross-domain policy file²: an XML blob that specifies the origins that are allowed to connect to that site [9]. The location of the policy file is itself subject to a number of restrictions, which make it more difficult for an attacker who has limited access to the target machine to generate a valid file. For instance, policy files hosted on ports ≥ 1024 cannot authorize access to ports < 1024 .

Flash Player uses the same general mechanism to control access both to raw sockets and to cross-domain HTTP requests. As with Java, Flash Player's consent mechanism was vulnerable to DNS rebinding attacks in the past³. Indeed, the mechanism described above where the cross-domain policy file is always checked is a response to some of these rebinding attacks which exploited a *time-of-check-time-of-use* (TOCTOU) issue between the browser's name resolution and that performed by Flash Player.

2) *JavaScript Cross-Origin Communication*: Until recently, network access for JavaScript applications was limited to making HTTP requests via `XMLHttpRequest`. Browsers heavily restrict these requests and forbid requesting cross-origin URLs [10]. Recently, browser vendors have added two mechanisms to allow web applications to escape (hopefully safely) from these restrictions.

a) *CORS: Cross-Origin Resource Sharing* (CORS) [11] allows web applications to issue HTTP requests to sites outside their origin. When a web application issues a cross-origin `XMLHttpRequest`, the browser includes the application's origin in the request in the `Origin` header. The server can authorize the application to read back the response by echoing the contents of the `Origin` request header in the `Access-Control-Allow-Origin` response header. This consent-based relaxation of the same-origin policy makes it easier for different web applications to communicate in the browser.

b) *WebSockets*: Although CORS is targeted only at HTTP requests, `WebSockets` [4] lets web applications open a socket connection to any server (whether or not the server is in the application's own origin) and send arbitrary data. This feature

²This description is a simplification of Flash Player's security policy [8].

³The DNS rebinding issues in Flash Player were fixed in version 9.0.115.0

is extremely useful, especially as an optimization for scenarios in which the server wishes to asynchronously send data to the client. Currently, such applications use a rather clumsy set of mechanisms generally known as Comet [12]. Like Flash Player and CORS, `WebSockets` uses a verified-origin mechanism to let the target server consent to the connection. Unlike Flash Player and CORS, the verification is performed over the same socket connection as will be used for the data (using a cryptographic handshake where the server replies to a client-provided nonce). This handshake is initiated by the browser and only after the handshake has completed does the browser allow the application to send data over the raw socket, which we further discuss in Section IV.

III. ATTACKS ON JAVA AND FLASH SOCKETS

As the history of DNS rebinding issues suggest, designing a robust same-origin or verified-origin policy is a challenging problem. Previous designs have been extremely subject to TOCTOU issues. In this section, we describe and demonstrate a new class of vulnerabilities which affect all the major existing and proposed same-origin and verified-origin policies, with the exception of CORS. Using the raw socket APIs available to web applications, our attacks exploit the existence of transparent proxies in networks and, in particular, their confusion about how to handle mismatches between the `HTTP Host` header and the destination IP address of the connection they are intercepting.

A. Vulnerabilities

Consider the situation in which the user is behind a transparent proxy and visits `attacker.com`. The attacker embeds a malicious SWF served from `attacker.com`, and the browser uses Flash Player to run the SWF. The attacker can now mount a number of different attacks, depending on how the proxy behaves.

1) *Route by Host Header*: When using a traditional proxy, the browser connects directly to the proxy and sends an HTTP request, which indicates to the proxy which resource the browser wishes to retrieve. When a transparent proxy intercepts an HTTP request made by a browser, the proxy has two options for how to route the request:

- The `HTTP Host` header.
- The IP address to which the browser originally sent the request.

Unfortunately, as described by Auger [2], if the proxy routes the request based on the `Host` header, an attacker can trick the proxy into routing the request to any host accessible to the proxy, as depicted in Figure 1:

- 1) The attacker hosts a permissive Flash socket policy server on `attacker.com:843` that allows access to every port from every origin.
- 2) The attacker's SWF requests to open a raw socket connection to `attacker.com:80` (which has IP address `2.2.2.2`).
- 3) Flash Player connects to `attacker.com:843` and retrieves the attacker's socket policy file, which indicates that the server has opted into the socket connection.

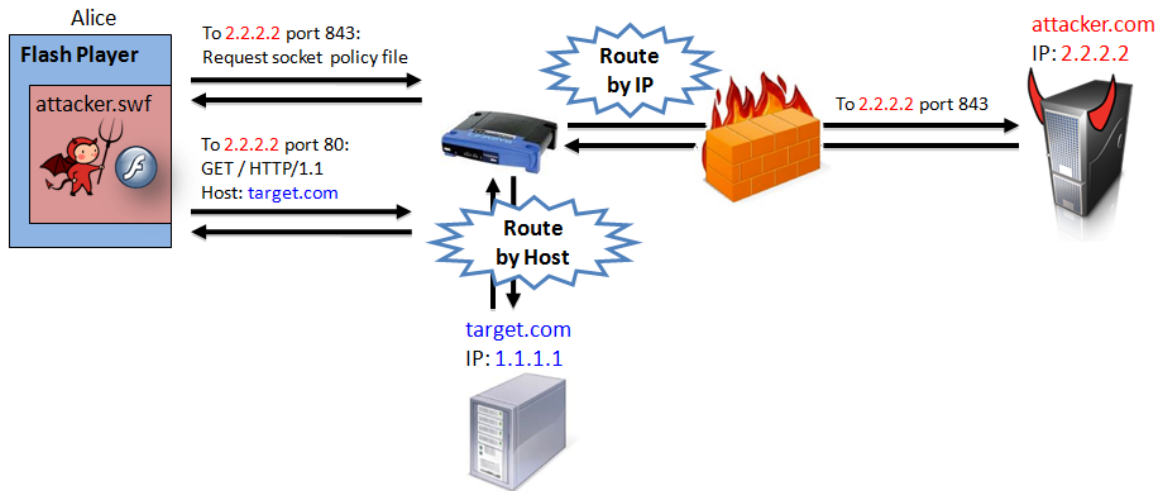


Fig. 1. IP hijacking attack

- 4) Flash Player lets the attacker's SWF open a new socket connection to `attacker.com:80`.
- 5) The attacker's SWF sends a sequence of bytes over the socket crafted with a fake Host header as follows:


```
GET / HTTP/1.1
Host: target.com
```
- 6) The transparent proxy treats these bytes as an HTTP request and routes the request according to the Host header (and *not* on the original destination IP address). Notice that the request is routed to `target.com:80` (which has an IP address of `1.1.1.1`).
- 7) The target server responds with the document for the URL `http://target.com/`, requested from the client's IP address, and the transparent proxy forwards the response to the attacker's SWF.

Notice that Flash Player authorized the attacker's SWF to open a socket to the attacker's server based on a policy file it retrieved from the attacker's server. However, the transparent proxy routed the request to a different server because the socket API let the attacker break the browser's security invariant that the Host header matched the destination IP address, leading to the vulnerability. Alternatively, the attacker can try to trick the proxy into tunneling a raw socket connection to the target server by using the HTTP CONNECT method [13] in step 5:

```
CONNECT target.com:80 HTTP/1.1
Host: target.com:80
```

By leveraging the user's machine to connect to other hosts in the Internet over these proxies, the attacker may hijack a user's IP address to perform misdeeds and frame the user. For example, the attacker may generate fake clicks on pay-per-click web advertisements to increase their advertising revenue [14], using different client IP addresses. IP hijacking attacks may also allow web attackers to access protected web sites that authenticate by IP address, or send spam email from the victim user's IP address.

An attacker can also exploit Java sockets in the same way. The attack steps are identical, except that the attacker need not host a policy file because Java implicitly grants applet the authority to open socket connections back to its origin server without requiring the server to consent.

2) *Cache by Host Header:* In the attacks described in the previous section, we considered transparent proxies that route HTTP requests according to the Host header. However, not all proxies are configured that way. Some proxies route the request to the original destination IP address, regardless of the Host header. Although these proxies are immune to IP hijacking attacks, we find that the attacker can still leverage some of these proxies to mount other attacks.

In particular, some transparent proxies that route by IP are also caching proxies. As with routing, proxies can cache responses either according to the Host header or according to the destination IP address. If a proxy routes by IP but caches according to the Host header, we discover that the attacker can instruct the proxy to cache a malicious response for an arbitrary URL of the attacker's choice, as shown in Figure 2:

- 1) The attacker's Java applet opens a raw socket connection to `attacker.com:80` (as before, the attacker can also a SWF to mount a similar attack by hosting an appropriate policy file to authorize this request).
- 2) The attacker's Java applet sends a sequence of bytes over the socket crafted with a forged Host header as follows:


```
GET /script.js HTTP/1.1
Host: target.com
```
- 3) The transparent proxy treats the sequence of bytes as an HTTP request and routes the request based on the original destination IP, that is to the attacker's server.
- 4) The attacker's server replies with malicious script file with an HTTP Expires header far in the future (to instruct the proxy to cache the response for as long as possible).
- 5) Because the proxy caches based on the Host header, the proxy stores the malicious script file in its

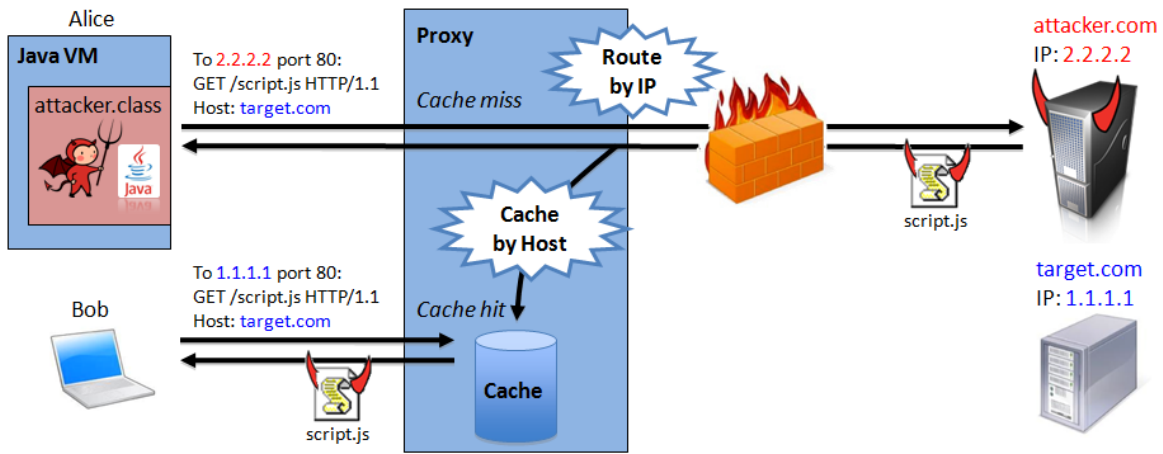


Fig. 2. Cache poisoning attack

cache as `http://target.com/script.js`, not as `http://attacker.com/script.js`.

- 6) In the future, whenever any client requests `http://target.com/script.js` via the proxy, the proxy will serve the cached copy of the malicious script.

One particularly problematic variant of this attack is for the attacker to poison the cache entry for Google Analytics, `http://www.google-analytics.com/ga.js`. Every user of the proxy (possibly the entire enterprise) will now load the attacker's malicious JavaScript into every page that uses Google Analytics, which is approximately 57% of the top 10,000 web sites [15]. Because the Google Analytics JavaScript runs with the privileges of the embedding web site, the attacker is able to effectively mount a persistent cross-site scripting attack against the majority of the Internet, as viewed by users of the proxy.

B. Experiment

The attacks described above have very specific network configuration requirements. To determine how commonplace these network configurations are on the Internet, we developed proof-of-concept exploits for both the IP hijacking and cache poisoning attacks using both Flash Player and Java. We then ran an advertisement on a public advertising network that mounted the attacks against servers in our laboratory.

- 1) *Methodology*: Our experiment consisted of two machines in our laboratory, with different host names and IP addresses. One machine played the role of the target server and the other played the role of the attacking server. The target was a standard Apache web server. The attacking server ran a standard Apache web server and a Flash socket policy server on port 843. We used a rich media banner advertisement campaign on an advertising network to serve our experimental code to users across the world. Our advertisement required no user interaction, and was designed to perform the following tasks in the user's web browser:

- a) *IP Hijacking*: Our advertisement opens a raw socket connection back to the attacking server using both Java and Flash Player. The attacking server runs a custom Flash socket policy server on port 843 that allows Flash socket connections to port 80 from any origin. Upon a successful connection, the advertisement spoofs an HTTP request over the socket by sending the following request:

```
GET /script.php/<random> HTTP/1.1
Host: target.com
```

The attacking server and the target server each host a PHP file at `/script.php`, but because these files are different we can easily determine which server the request went to. The random value on the end of the URL serves to bypass caches used by plug-ins, browsers, or proxies. Alternatively, we could have included the random value in the query string (i.e., after a `?` character) but some caching proxies treat URLs containing query strings inconsistently.

If the HTTP response was from the target server instead of from the attacking server, that is direct evidence that the request was routed by the `Host` header, which implies that the user is vulnerable to IP hijacking.

- b) *Cache Poisoning*: In the previous test, the script files were served with `Cache-Control: public`, `Last-Modified` and `Expires` response headers that allowed them to be cached for one year. To check whether the socket connection has poisoned the proxy's cache, we added a script tag to our advertisement that attempts to load a script from the target server at `http://target.com/script.php/<random>`, reusing the random value from the previous request.

Because the random value was only used previously via the socket API, this URL will not be present in the browser's HTTP cache (as the browser does not observe the bytes sent over the socket). By checking the contents of the response (specifically, a JavaScript variable), we can determine whether the script was from the attacker or the target server. If we

| | Flash Player | Java |
|--|--------------|-------------|
| Spoof request routed to target* | 3152 | 2109 |
| Spoof request routed to attacker | 47839 | 26759 |
| Script file cached from target | 51163 | 26612 |
| Script file cached from attacker† | 108 | 53 |

TABLE I
HTTP HOST HEADER SPOOFING VIA PLUG-IN SOCKETS

| | POST-based | Upgrade-based | CONNECT-based |
|--|-------------|---------------|---------------|
| Handshake pass and spoof request ignored | 47741 | 47162 | 47204 |
| Spoof request routed to target* | 1376 | 1 | 0 |
| Spoof request routed to attacker | 97 | 174 | 2 |
| Script file cached from target | 54519 | 54526 | 54534 |
| Script file cached from attacker† | 15 | 8 | 0 |

TABLE II
HTTP HOST HEADER SPOOFING VIA HTML5 WEBSOCKET STRAWMAN PROTOCOLS

* Allows attacker to open a direct socket from the client to an arbitrary server

† Allows attacker to poison the HTTP cache of all clients of the proxy

receive the version of the script hosted on the attack server, we can deduce that a transparent proxy has cached the response.

2) *Results*: We ran our advertisement on five successive days in March 2011, spending \$100 in total. We garnered a total of 174,250 unique impressions. We discarded repeat visits by the same users by setting a cookie in the user’s browser. The advertisement ran our JavaScript, SWF, and Java bytecode without user intervention and sent results back to server in our laboratory after completing the experiment. If the user closed the browser window or navigated away before the experiment finished running, we did not receive the results from that part of the experiment. We collected 51,273 results from SWFs and 30,045 results from Java applets (19,117 of the impressions produced results from both tests). The most likely reason for the low response rate is that the loading time of our SWF and Java applet was noticeably slow, and users did not stay on the page long enough for the experiment to run. Our experimental results show that both IP hijacking attacks and cache poisoning exist in real world scenarios, as shown in Table I.

a) *IP Hijacking*: In the IP hijacking test using Flash sockets, we observed that the spoofed request was routed back to the attacking server on 47,839 of 51,273 impressions (93.3%), suggesting that the client made a direct connection or the network intermediaries routed regardless of the `Host` header. We logged 233 of 51,273 impressions (0.4%) where the Flash socket failed to open, possibly due to firewalls that blocked port 843, preventing Flash Player from fetching the socket policy file. There were 49 cases where the client received an HTML error message, possibly generated by a transparent proxy that blocked the spoofed request. On 3,152 impressions (6.1%) the spoofed request was routed by the `Host` header to the target server, indicating vulnerability to IP hijacking.

Using Java sockets, we observed that 26,759 of 30,045 impressions (89.1%) received the response from the attacker’s server, implying that they were routing on IP. Out of 30,045 impressions, there were 1,134 (3.8%) connection errors that threw Java exceptions and 43 that received an HTML error

message. We found that 2,109 of 30,045 impressions (7%) routed on the `Host` header, allowing IP hijacking attacks.

b) *Cache Poisoning*: In the cache poisoning test using Flash sockets, we observed that 51,163 of 51,273 impressions (99.8%) were able to fetch the script from the target. There were 2 cases where the client reported an error response. However, we discovered that the cache poisoning attack was successful on 108 of 51,273 impressions (0.21%). This suggests that some transparent proxies route HTTP requests by IP but cache according to the `Host` header.

In our cache poisoning test using Java sockets, we observed 26,612 of 30,045 impressions (88.6%) retrieved the response from the target server. We observed that 3,680 of 30,045 impressions (12.2%) caused exceptions when using Java to interrogate the results of the second query, which we were unable to determine whether the cache poisoning succeeded or not. Similarly to the results using Flash sockets, there were 53 of 30,045 impressions (0.18%) that reported a successful cache poisoning attack.

Our results show that the attacker may achieve a cost efficiency of 1.08 successful cache poisoning attacks per dollar spent, using Flash sockets on advertising networks. Note that each successful cache poisoning attack would in effect compromise other users of the vulnerable proxy, beyond our measurement.

IV. ATTACKS ON WEBSOCKET PROTOCOLS

One diagnosis of the cause of the Java and Flash socket vulnerabilities is that both use an out-of-band mechanism to authorize socket connections. Because intermediaries are oblivious to these out-of-band signals, they misinterpret the information sent over the socket by the attacker. In this section, we consider three in-band signaling mechanisms for authorizing socket connections, all based on HTTP. The first is a `POST`-based handshake of our own invention to illustrate some of the design issues. The second is the state-of-the-art `Upgrade`-based handshake used by HTML5. The third is an experimental

CONNECT-based handshake that we designed in attempt to prevent attacks.

A. POST-based Handshake

1) *Design*: One natural approach to designing an in-band signaling mechanism is to model the handshake after HTTP. The idea here is that until we have established the server's consent to receive WebSockets traffic, we will not send any data that the attacker could not already have generated with existing browser functionality—with the HTML form element being the most powerful piece of syntax in this respect—so what could possibly go wrong? This should protect servers which do not want to speak WebSockets from being sent WebSockets data. With this goal in mind, consider the following strawman handshake based on an HTTP POST request:

Client → *Server*:

```
POST /path/of/attackers/choice HTTP/1.1
Host: host-of-attackers-choice.com
Sec-WebSocket-Key: <connection-key>
```

Server → *Client*:

```
HTTP/1.1 200 OK
Sec-WebSocket-Accept: <connection-key>
```

By echoing the connection key to the client, the server consents that it accepts the WebSocket protocol. If WebSockets are less generative than the form element, then we might believe that adding WebSockets support to browsers does not increase the attack surface.

2) *Vulnerabilities*: Unfortunately, using this handshake WebSockets are not less generative than the HTML form element. For example, WebSocket applications can generate data that appear as framing escapes and confuse network intermediaries into handling subsequent data as new HTTP connections, instead of a continuous single HTTP connection expressed by the form element. Although we have accomplished our initial goal of not sending any non-HTTP data to WebSockets servers, we can still confuse transparent proxies.

Consider an intermediary examining packets exchanged between the browser and the attacker's server. As above, the client requests WebSockets and the server agrees. At this point, the client can send any traffic it wants on the channel. Unfortunately, the intermediary does not know about WebSockets, so the initial WebSockets handshake just looks like a standard HTTP request/response pair, with the request being terminated, as usual, by an empty line. Thus, the client program can inject new data which looks like an HTTP request and the proxy may treat it as such. So, for instance, he might inject the following sequence of bytes:

```
GET /sensitive-document HTTP/1.1
Host: target.com
```

When the intermediary examines these bytes, it might conclude that these bytes represent a second HTTP request over the same socket. If the intermediary is a transparent proxy, the

intermediary might route the request or cache the response according to the forged Host header, discussed in Section III.

3) *Experiment*: To evaluate the practicality of mounting IP hijacking and cache poisoning attacks with the WebSocket handshakes, we implemented prototypes for each WebSocket handshake using Flash sockets and a WebSocket server written in Python. We reused the system from the Java and Flash socket experiment with the following changes. We setup a custom multiplexing server at port 80 on the attacking server, which forwards requests to either a standard Apache server or the WebSocket server depending on the request headers. We ran an advertisement campaign for four successive days in November 2010, spending \$20 in the Philippines and \$80 globally. Our advertisement contains a SWF which performs the WebSocket handshake, spoofs an HTTP request upon handshake success, and instructs the browser to request a script from the target server using a script tag. We experimented with how intermediaries process each WebSocket handshake. Table II shows our results.

Out of a total of 54,534 impressions, 49,218 (90.2%) succeeded with the POST-based handshake and 5,316 (9.4%) failed. Out of the 49,218 impressions on which we were able to run our IP hijacking test, 47,741 (96.9%) reported that no intermediaries were confused when sending the spoofed HTTP request. However, we found that the IP hijacking attack succeeded on 1,376 of 49,218 impressions (2.8%), where the client was behind a Host-routing proxy. There were 97 of 49,218 impressions (0.2%) where the spoofed request was routed by IP and 4 that received an HTML error. We ran the cache poisoning test on the clients that succeeded with the POST-based handshake, and found 15 successful cache poisoning attacks. These results show that the POST-based handshake is vulnerable to both attacks.

B. Upgrade-based Handshake

1) *Design*: In an attempt to improve the security of its socket handshake, HTML5 uses HTTP's Upgrade mechanism to upgrade from the HTTP protocol to the WebSocket protocol. HTTP's Upgrade mechanism is a generic mechanism for negotiating protocols using HTTP which was originally designed for layering TLS over HTTP. HTTP's Upgrade mechanism has two pieces: a Connection header whose value is the string "Upgrade" and an Upgrade header whose value is the name of the protocol to which the client wishes to switch. Below is a simplified version of the HTML5 WebSocket handshake using HTTP's Upgrade mechanism.

Client → *Server*:

```
GET /path/of/attackers/choice HTTP/1.1
Host: host-of-attackers-choice.com
Connection: Upgrade
Sec-WebSocket-Key: <connection-key>
Upgrade: WebSocket
```

Server → *Client*:

```
HTTP/1.1 101 Switching Protocols
```

```
Connection: Upgrade
Upgrade: WebSocket
Sec-WebSocket-Accept:
  HMAC(<connection-key>, "...")
```

2) *Vulnerabilities*: Unfortunately, HTTP's Upgrade mechanism is virtually unused in practice. Instead of layering TLS over HTTP using Upgrade, nearly every deployment of HTTP over TLS uses a separate port, typically port 443 (the generic name for this mode is HTTPS [16]). Consequently, many organizations are likely to deploy network intermediaries that fail to implement the Upgrade mechanism because these intermediaries will largely function correctly on the Internet today. Implementers and users of these intermediaries have little incentive to implement Upgrade, and might, in fact, be unaware that they do not implement the mechanism.

To an intermediary that does not understand HTTP's Upgrade mechanism, the HTML5 WebSocket handshake appears quite similar to our strawman POST-based handshake. These intermediaries are likely to process the connection the same way for both the POST-based handshake and the Upgrade-based handshake. If such an intermediary is vulnerable to the attacks on the POST-based handshake, the intermediary is likely to be vulnerable to the same attacks when using the Upgrade-based handshake.

3) *Experiment*: In our experiment, we tested how intermediaries in the wild process the Upgrade-based handshake. Out of a total of 54,534 impressions, 47,338 (86.8%) succeeded with the handshake and 7,196 (13.2%) failed. The handshake failed more often than the POST-based handshake, possibly when the Upgrade mechanism was unsupported and, perhaps, stripped. Out of the 47,338 impressions on which we were able to run our IP hijacking test, 47,162 (99.6%) did not receive a response after spoofing an HTTP request. We noticed that the IP hijacking attack succeeded on 1 impression, where the client was behind a Host-routing proxy. There were 174 of 47,338 impressions (0.37%) where the spoofed request was routed by IP. One impression received an HTML error message.

Out of the 47,338 impressions that succeeded the Upgrade-based handshake, we ran the cache poisoning test and found 8 successful cache poisoning attacks. The 8 impressions were also vulnerable to cache poisoning when using the POST-based handshake.

C. CONNECT-based Handshake

1) *Design*: Rather than relying upon the rarely used HTTP Upgrade mechanism to inform network intermediaries that the remainder of the socket is not HTTP, we consider using HTTP's CONNECT mechanism. Because CONNECT is commonly used to establish opaque tunnels pass TLS traffic, transparent proxies are likely to interpret this request as an HTTPS connect request, assume the remainder of the socket is unintelligible, and simply route all traffic transparently based on the IP. We create a strawman handshake based on the CONNECT mechanism.

Client → Server:

```
CONNECT websocket.invalid:443 HTTP/1.1
```

```
Host: websocket.invalid:443
Sec-WebSocket-Key: <connection-key>
Sec-WebSocket-Metadata: <metadata>
```

Server → Client:

```
HTTP/1.1 200 OK
Sec-WebSocket-Accept: <hmac>
```

where <connection-key> is a 128-bit random number encoded in base64 and <metadata> is various metadata about the connection (such as the URL to which the client wishes to open a WebSocket connection). In the server's response, <hmac> is the HMAC of the globally unique identifier 258EAF5-E914-47DA-95CA-C5AB0DC85B11 under the key <connection-key> (encoded in base64). By sending the <hmac> value, the server demonstrates to the client that it understands and is willing to speak the WebSocket protocol because computing the <hmac> value require "knowledge" of an identifier that is globally unique to the WebSocket protocol.

Notice that instead of using the destination server's host name, we use an invalid host name (per RFC 2606 [17]). Any intermediaries that do not recognize the WebSocket protocol but understand this message according to its HTTP semantics will route the request to a non-existent host and fail the request.

2) *Experiment*: We tested whether the CONNECT-based handshake would resist transparent proxy attacks in the real world. Out of a total of 54,534 impressions, 47,206 (86.6%) succeeded with the handshake and 7,328 (13.4%) failed. Out of the 47,206 impressions on which we were able to run our IP hijacking test, only three did receive a response after spoofing an HTTP request. We observed that the IP hijacking attack did not succeed on any clients. We logged 1 impression that returned an HTML error message. We observed 2 impressions where the spoof request was routed by IP to the attacking server, however none indicated proxy routing based on the Host header. It appears that these proxies simply passed the CONNECT to our server untouched and then treated the next spoofed request as if it were a separate request routed by IP. We proceeded to the cache poisoning test and did not find successful cache poisoning attacks.

D. Our Proposal

1) *Design*: In our experiments, we found successful attacks against both the POST-based handshake and the Upgrade-based handshake. For the CONNECT-based handshake, we observed two proxies which appear not to understand CONNECT but simply to treat the request as an ordinary request and then separately route subsequent requests, with all routing based on IP address. Although these proxies did not cache, it is possible that proxies of this type which cache do exist—though our data suggest that they would be quite rare. In this case the attacker would be able to mount a cache poisoning attack.

A mitigation for these attacks is to mask all the attacker-controlled bytes in the raw socket data with a stream cipher. The stream cipher is not to provide confidentiality from

eavesdroppers but to ensure that the bytes on the wire appear to be chosen uniformly at random to network entities that do not understand the WebSocket protocol, making it difficult for the attacker to confuse the receiver into performing some undesirable action.

We propose masking the metadata in the initial handshake and all subsequent data frames with a stream cipher, such as AES-128-CTR. To key the encryption, the client uses HMAC of the globally unique identifier C1BA787A-0556-49F3-B6AE-32E5376F992B with the key `<connection-key>`. However, encrypting the raw socket writes as one long stream is insufficient because the attacker learns the encryption key in the handshake thus can generate inputs to the socket write function that produce ciphertexts of his choice. Instead, we encrypt each protocol frame separately, using a per-frame random nonce as the top part of the CTR counter block, with the lower part being reserved for the block counter. From the perspective of the attacker, this effectively randomizes the data sent on the wire even if the attacker knows the key exchanged in the handshake. Note that each protocol frame must be encrypted with a fresh nonce and that the browser must not send any bytes on the wire until the browser receives the entire data block from the application. Otherwise, the attacker could learn the nonce and adjust the rest of the input data based on that information.⁴ This mitigation comes at a modest performance cost and some cost in packet expansion for the nonce, which needs to be large enough that the attacker’s chance of guess the nonce is sufficiently low.

In the case that the cost of encryption is a burden, Stachowiak [19] suggests using a simple XOR cipher as a lightweight alternative to using AES-128-CTR. In particular, the client generates a fresh 32 bit random nonce for every frame, and the plaintext is XORed with a pad consisting of the nonce repeated. Because the nonce is unknown to the attacker prior to receiving the corresponding data frame, the attacker is unable to select individual bytes on the wire. However, because the pad repeats, the attacker is able to select correlations between the bytes on the wire, but we are unaware of how to leverage that ability in an attack.

Other proposals with simpler transformations have been discussed in the WebSocket protocol working group, such as flipping the first bit in the frame, or escaping ASCII characters and carriage returns in the handshake. However, these proposals do not protect servers or intermediaries with poor implementation that skip non-ASCII characters. Moreover, using cryptographic masking also mitigates other attack vectors, such as non-HTTP servers that speak protocols with non-ASCII bytes. We believe masking is a more robust solution to these attacks that is more likely to withstand further security analysis.

2) *Experiment*: We evaluated the network performance of WebSockets using no masking, XOR masking (with 32 bit nonces) and AES-128-CTR masking (with 32, 64 and 128 bit nonces), modified on a Java implementation [20]. From

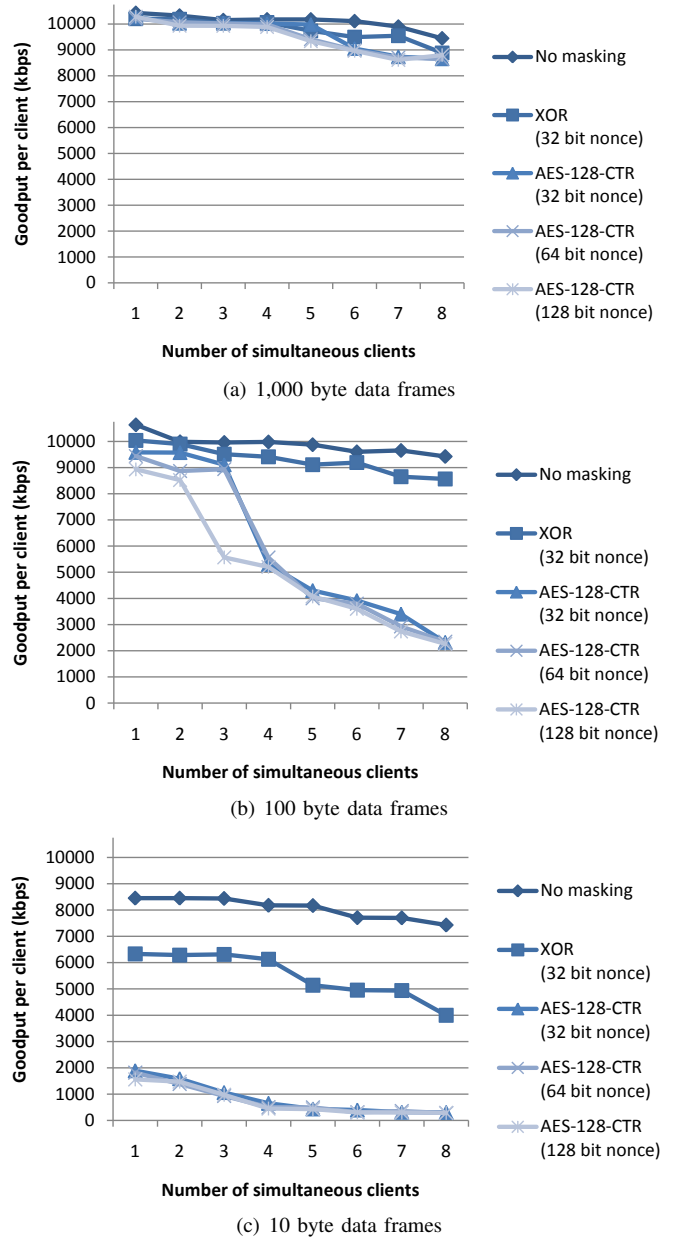


Fig. 3. Performance of WebSocket data frames

slicehost.com, we acquired a 1,024 MB RAM machine as the server with uncapped incoming bandwidth and eight 256 MB RAM machines as the clients, each with 10 Mbps outgoing bandwidth. In our evaluation, we measured the elapsed time for each client to send 10 MB of application data to the server with various frame sizes, while the server handles up to 8 clients simultaneously. Results for sending 1,000 byte data frames, 100 byte data frames and 10 byte data frames are shown in Figure 3(a), Figure 3(b) and Figure 3(c), respectively. We observe that AES-128-CTR masking induces little overhead when the data frame size is as large as 1,000 bytes. However, the performance of AES-128-CTR masking drops off significantly for smaller data frames in comparison

⁴A similar condition applies to TLS [18] packet writes.

with no masking, whereas XOR masking still performs at acceptable speeds.

3) *Adoption*: We reported the vulnerabilities to the IETF WebSocket protocol working group in November 2010. Due to concerns about these attacks, Firefox [21] and Opera [22] temporarily disabled the WebSocket protocol. In response to our suggestion, the working group reached consensus to prevent the attacker from controlling the bytes sent on the wire by requiring XOR-based masking. Internet Explorer adopted frame masking in their WebSocket prototype using Silverlight plug-in in HTML5 Labs [23]. We hope to assist the Flash Player and Java plug-ins in addressing these issues in the near future.

4) *Discussion*: In our study, we observe a number of misbehaving network intermediaries. Unfortunately, we are unable to determine which specific proxy implementations are vulnerable because the misbehaving proxies were almost entirely transparent. For example, the proxies did not announce their presence using the HTTP `Via` header, as required by the HTTP specification. Moreover, the vulnerable behavior might actually be the result of a chain of proxies, none of which are individually vulnerable.

One approach to resolving these vulnerabilities is to wait for misbehaving proxies to be replaced. However, the time horizon for replacing these proxies is unbounded. Rather than wait for these proxies to be fixed, we recommend that browser vendors resolve the issue in the HTML5 WebSocket protocol itself, as they have done. Further, we recommend that the appropriate vendors fix the related vulnerabilities in Flash Player and Java. (Note: users and enterprises can mitigate silent exploitation of these plug-in vulnerabilities by disabling the plug-ins by default and using a “click-to-Flash” authorization model.)

V. RELATED WORK

A. Cross-Protocol Attacks

Cross-protocol attacks are used to confuse a server or an intermediary into associating a request with an incorrect protocol. We described an instance of a cross protocol attack between HTTP and the WebSocket protocol. Topf [24] describes an attack that uses HTML forms to send commands to servers running ASCII based protocols like SMTP, NNTP, POP3, IMAP, and IRC. To prevent these attacks, browsers restrict access to well-known ports of vulnerable applications, such as port 25 for SMTP. This defense cannot be applied to WebSockets because WebSockets operates over port 80, the same port as with HTTP, for compatibility. We suspect there are other forms of cross-protocol attacks and expect to address more of these problems in future work.

B. HTTP Cache Poisoning

Bueno [25] describes an HTTP cache poisoning attack on web pages that rely on the value of the HTTP `Host` header to generate HTML links. In particular, a malicious client sends an HTTP request with a crafted `Host` header, causing the server to rewrite links with an arbitrary string provided by the attacker. If there is any caching going on by proxies along the way, other clients will get the exploited page with injected text. A

mitigation for these attacks is to not generate any page content using the `Host` header. In comparison, our cache poisoning attacks do not rely on the usage of `Host` header in the target page, and allow the attacker to poison the proxy’s cache for an arbitrary URL on any target host.

C. HTTP Response Splitting

In an HTTP response splitting attack [26], the attacker sends a single HTTP request that tricks the benign server into generating an HTTP response that is misinterpreted by the browser or an intermediary as two HTTP responses. Typically, the malicious request contains CRLF sequences that are reflected by the server into the output stream and appear to terminate the first response, letting the attacker craft the byte sequence that the browser or intermediary interprets as the second response. The attacker can mount a cache poisoning attack by sending a second request to a benign server, which causes the browser or proxy associates with the second “response” and stores in its cache. Servers can prevent the attack by sanitizing data and not allowing CRLF in HTTP response headers. In our work, we introduce new cache poisoning attacks against transparent proxies, which are not addressed by previous mitigations.

D. Pretty-Bad-Proxy

Chen et. al. [27] introduce a series of attacks in which a malicious proxy breaks the end-to-end security guarantees of the TLS protocol by injecting messages that are interpreted as HTTPS responses by the browser. A malicious proxy can trick browsers into running a script of the attacker’s choice in the security context of a target server by embedding scripts in HTTP error messages or by redirecting script requests to malicious servers using HTTP redirects. Browsers have mitigated these vulnerabilities by ignoring the proxy redirection and error messages received prior to completing the TLS handshake. Our work does not focus on malicious proxies. Rather, we are interested in benign-but-confused proxies. However, one should always be wary of malicious proxies when designing a secure communication protocol.

VI. CONCLUSION

Although raw socket access is an important capability for full-featured browser-based applications, providing sockets safely has proven to be challenging. Although raw socket access requires the destination server’s consent to receive raw socket traffic, our results demonstrate that raw sockets can still be abused in the presence of certain transparent proxies. Our experiments show that approximately 7% of browsers are behind proxies with implementation errors that may enable attack via one of these vectors.

The designers of consent protocols should consider how the attacker can manipulate these protocols to exploit network intermediaries that unintentionally proxy the consent of the remote server without understanding its semantics. We propose improving the security of current consent mechanisms by encrypting all the attacker-controlled bytes sent over the

wire using per-frame random nonces so that raw socket traffic appears random to oblivious network intermediaries. In response to our suggestion, the WebSocket protocol working group has introduced frame masking, improving the security of WebSockets.

REFERENCES

- [1] C. Jackson, A. Barth, A. Bortz, W. Shao, and D. Boneh, "Protecting browsers from dns rebinding attacks," in *Proceedings of the 14th ACM Conference on Computer and Communications Security*, 2007.
- [2] R. Auger, "Socket capable browser plugins result in transparent proxy abuse," 2010, http://www.thesecuritypractice.com/the_security_practice/TransparentProxyAbuse.pdf.
- [3] I. Fette, "The WebSocket protocol," 2011, <http://tools.ietf.org/html/draft-ietf-hybi-thewebsocketprotocol>.
- [4] I. Hickson, "The Web Sockets API," 2009, <http://www.w3.org/TR/websockets/>.
- [5] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, "Hypertext Transfer Protocol – HTTP/1.1," RFC 2616 (Draft Standard), Internet Engineering Task Force, Jun. 1999, updated by RFCs 2817, 5785. [Online]. Available: <http://www.ietf.org/rfc/rfc2616.txt>
- [6] D. D. Edward, E. W. Felten, and D. S. Wallach, "Java security: From hotjava to netscape and beyond," in *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, 1996.
- [7] H. Wang, X. Fan, J. Howell, and C. Jackson, "Protection and communication abstractions for web browsers in mashups," in *Proceedings of 21st ACM SIGOPS Symposium on Operating Systems Principles (SOSP)*, 2007.
- [8] Adobe, "White paper: Adobe flash player 10 security," 2008, http://www.adobe.com/devnet/flashplayer/articles/flash_player10_security_wp.html.
- [9] Adobe, "Cross-domain policy file specification," 2010, http://www.adobe.com/devnet/articles/crossdomain_policy_file_spec.html.
- [10] M. Zalewski, "Browser security handbook," <http://code.google.com/p/browsersec/wiki/Main>.
- [11] A. van Kesteren, "Cross-Origin Resource Sharing," 2010, <http://www.w3.org/TR/cors/>.
- [12] A. Russell, "Comet: Low Latency Data for the Browser," 2006, <http://infrequently.org/2006/03/comet-low-latency-data-for-the-browser/>.
- [13] R. Khare and S. Lawrence, "Upgrading to TLS Within HTTP/1.1," RFC 2817 (Proposed Standard), Internet Engineering Task Force, May 2000. [Online]. Available: <http://www.ietf.org/rfc/rfc2817.txt>
- [14] V. Anupam, A. Mayer, K. N. an Benny Pinkas, and M. K. Reiter, "On the security of pay-per-click and other web advertising schemes," in *Proceedings of the 8th International Conference on World Wide Web*, 1999.
- [15] BuiltWith, "Google Analytics Usage Statistics," 2011, <http://trends.builtwith.com/analytics/Google-Analytics>.
- [16] E. Rescorla, "HTTP Over TLS," RFC 2818 (Informational), Internet Engineering Task Force, May 2000, updated by RFC 5785. [Online]. Available: <http://www.ietf.org/rfc/rfc2818.txt>
- [17] D. Eastlake 3rd and A. Panitz, "Reserved Top Level DNS Names," RFC 2606 (Best Current Practice), Internet Engineering Task Force, Jun. 1999. [Online]. Available: <http://www.ietf.org/rfc/rfc2606.txt>
- [18] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2," RFC 5246 (Proposed Standard), Internet Engineering Task Force, Aug. 2008, updated by RFCs 5746, 5878. [Online]. Available: <http://www.ietf.org/rfc/rfc5246.txt>
- [19] M. Stachowiak, "Re: [hybi] handshake was: The websocket protocol issues." 2010, <http://www.ietf.org/mail-archive/web/hybi/current/msg04379.html>.
- [20] J. Tamplin, "Sample code for evaluation of WebSocket draft proposals," 2011, <http://code.google.com/p/websocket-draft-eval/>.
- [21] C. Heilmann, "WebSocket disabled in Firefox 4," 2010, <http://hacks.mozilla.org/2010/12/websockets-disabled-in-firefox-4/>.
- [22] A. van Kesteren, "Disabling the WebSocket Protocol," 2010, <http://annevankesteren.nl/2010/12/websocket-protocol-vulnerability>.
- [23] C. Caldato, "The Updated WebSockets Prototype," 2011, <http://blogs.msdn.com/b/interoperability/archive/2011/02/09/the-updated-websockets-prototype.aspx>.
- [24] J. Topf, "Html form protocol attack," 2001, <http://www.remote.org/jochen/sec/hfpa/hfpa.pdf>.
- [25] C. Bueno, "HTTP Cache Poisoning via Host Header Injection," 2008, <http://carlos.bueno.org/2008/06/host-header-injection.html>.
- [26] A. Klein, "Divide and conquer - HTTP response splitting, web cache poisoning attacks, and related topics," 2004, http://packetstormsecurity.org/papers/general/whitepaper_httppresponse.pdf.
- [27] S. Chen, Z. Mao, Y.-M. Wang, and M. Zhang, "Pretty-bad-proxy: An overlooked adversary in browsers' https deployments," in *Proceedings of the 2009 30th IEEE Symposium on Security and Privacy*, 2009.