BrowserAudit:
Automated Testing of Browser Security Features

Charlie
Hothersall-Thomas
Netcraft Ltd, UK
me@charlie.ht

Sergio Maffeis
Department of Computing
Imperial College London, UK
maffeis@doc.ic.ac.uk

Chris Novakovic
Department of Computing
Imperial College London, UK
c.novakovic@imperial.ac.uk

ABSTRACT

The security of the client side of a web application relies on browser features such as cookies, the same-origin policy and HTTPS. As the client side grows increasingly powerful and sophisticated, browser vendors have stepped up their offering of security mechanisms which can be leveraged to protect it. These are often introduced experimentally and informally and, as adoption increases, gradually become standardised (e.g., CSP, CORS and HSTS). Considering the diverse landscape of browser vendors, releases, and customised versions for mobile and embedded devices, there is a compelling need for a systematic assessment of browser security.

We present BrowserAudit, a tool for testing that a deployed browser enforces the guarantees implied by the main standardised and experimental security mechanisms. It includes more than 400 fully-automated tests that exercise a broad range of security features, helping web users, application developers and security researchers to make an informed security assessment of a deployed browser. We validate BrowserAudit by discovering both fresh and known security-related bugs in major browsers.

Categories and Subject Descriptors

D.4.6 [**Operating systems**]: Security and Protection

Keywords

Web security, web browser testing, same-origin policy, Content Security Policy, Cross-Origin Resource Sharing, clickjacking, cookies

1. INTRODUCTION

Personal data, business transactions, critical infrastructure and even cars, refrigerators and lightbulbs are exposed through web interfaces to a wide variety of web browsers. Hence, the browser plays a key role in the modern information infrastructure, as the main gateway to access the information and capabilities made available online.

As such, browsers need to offer a variety of standardised security mechanisms which can be relied upon uniformly by the client side of web applications, in order to deliver security guarantees to their users. For example, the same-origin policy (SOP) [34] is effective at preventing a range of cross-site scripting (XSS) attacks [38] against users' web browsers and is an integral aspect of modern web-based security. On the other hand, it is sometimes excessively strict; for instance, it forbids the sharing of information between different subdomains, a common requirement of large web sites. It is also coarse-grained, and several attempts have been made to enforce finer-grained access control [41, 39] and origins [22, 23, 29] in the browser. A variety of contemporary web browsers implement the Cross-Origin Resource Sharing (CORS) [46] standard, which may be used to control the flow of information between server-side resources and client-side scripts that attempt to access those resources via APIs. However, even fully-compliant implementations of the SOP and CORS mechanisms in some cases do not regulate access to other resources, such as images, embedded objects and web fonts, that can leave web applications vulnerable to cross-site request forgery (CSRF) attacks [20], clickjacking [36], framebusting [43] and CSS-based attacks [33]. The Content Security Policy (CSP) standard [45] enables much finer-grained control over the loading of arbitrary resources on a web page, mitigating several of these issues. These are just some examples of established and emerging security mechanisms offered by modern browsers.

Such mechanisms are often introduced experimentally and informally. As adoption increases, they gradually become standardised, and after numerous security reviews and bug reports they can eventually be relied upon consistently across browsers [19, 37, 20]. Reaching that stage is not easy. For example, correctly implementing the CSP specification is non-trivial: it is a lengthy document with many cross-references to other standards and RFCs, many of which have been superseded by newer (and conflicting) standards and RFCs. It is possible that a browser vendor could incorrectly implement some part of the CSP and thus fail to provide some of its security guarantees to their users. There is therefore a need for an automated tool that enables browser developers to complement low-level unit tests targeted at isolated source code modules with high-level testing of the effectiveness of the implementation of the security features once the browser is deployed.

In this paper we introduce BrowserAudit, a framework for testing whether a deployed browser correctly enforces the security guarantees implied by the main standardised

security mechanisms. For practical purposes, we present BrowserAudit as a standalone web application that automatically tests the browser used to access it. BrowserAudit has been designed with different sets of users in mind. A casual web user can run the tests to gain a simple security assessment of their browser: critically vulnerable, non-critically vulnerable, or okay. With the recent surge of security breaches reported in the news, people are becoming increasingly security-conscious and we believe there is an increasing demand for tools that inform the public about security. A security researcher can benefit even more, viewing a detailed breakdown of each test result, and seeing which security features passed our tests and which had problems. We display textual descriptions for each category of tests and the client-side source code of the tests. Browser developers can use BrowserAudit to debug their security features and web developers can use it as a way to ascertain the security capability of users' browsers (Section 2). We chose to implement a careful selection of tests that covers both the most important browser security mechanisms that should be implemented in any browser, and some of the most promising experimental ones that are not yet widely implemented. Starting from the code of individual test cases, we identified and generalised common patterns in order to automatically generate hundreds of tests. BrowserAudit automatically tests over 400 behaviours where a certain action should either be allowed or blocked according to an implied browser security policy (Section 3).

We designed BrowserAudit to be efficient and scalable, and we evaluated its performance and its accuracy extensively by running it on a number of browsers and platforms. Using BrowserAudit, we have discovered several previously unknown security bugs in recent versions of Mozilla Firefox, which we have reported to the developers (Section 4.4).

Whilst there are well-understood methodologies for generating unit tests for a given code base, there is no general solution to the problem of testing the end-to-end security behaviour of a family of applications (in our case web browsers) that must respect precise interoperability constraints (web standards) but can widely differ in implementation architectures, languages and design. Hence, we faced a significant challenge when developing our tests, carrying out a substantial amount of practical experimentation, guided by the official RFCs, our formal and informal models of web security, and a substantial body of academic and practical research on browser and web security (surveyed in Section 5.1). Although we believe that BrowserAudit is unique in its focus and breadth, we were inspired by a number of related web applications described in Section 5.2.

Contributions. Summarising, our main contributions are:

- We analysed the specifications of HTML5, CSP, CORS and HTTP Strict Transport Security (HSTS), identifying the concrete security guarantees implied by the proposed mechanisms. This allowed us to formulate precise goals for security test cases.

- We built a suite of more than 400 browser security tests, which brings together a wealth of explicit and implicit knowledge of the guarantees afforded by modern browser security mechanisms. We made the tests available to the community by open-sourcing the BrowserAudit code base [32].

- We implemented the first fully-automated web application that comprehensively tests browser security features and provides detailed information to a variety of user bases.

- We used BrowserAudit to discover previously unknown vulnerabilities in a major web browser.

2. DESIGN OVERVIEW

The goals underlying the design of BrowserAudit are the following:

- *Wide coverage*: BrowserAudit should demonstrate that a wide range of browser security mechanisms can be tested automatically, reliably and efficiently. Complete test coverage of any such mechanism is not practically feasible, and beyond the scope of this project.[1]

- *Extensibility*: By its very nature, BrowserAudit will always be a work in progress. As the browser threat landscape evolves, more tests will be needed to cover new security mechanisms, or to extend the coverage of existing ones. Our design should ease the task of creating, debugging and integrating additional test cases.

- *Ease of use*: BrowserAudit should be easily accessible on any modern browser connected to the Internet, without the need to install additional software. It should require no interaction from the user, otherwise running hundreds of tests would be impractical. Moreover, relying on user interaction would prevent the desired aim of running the tests transparently in the background.

- *Broad audience*: Our design should support a diverse range of users. A report on the security effectively offered by a deployed browser should benefit browser developers, penetration testers, security researchers and web users.

- *Scalability*: Our design should be scalable on the server side. Several users may be testing their browser at the same time, and many security tests concern features that involve communicating with the server.

We now sketch the architecture of BrowserAudit and highlight the main design choices. We defer further implementation details to Sections 3 and 4.1.

2.1 User Experience

BrowserAudit is accessible by simply pointing the browser to be tested to `https://browseraudit.com/`. This is a landing page that briefly describes the aims of the project and contains a "Test me" button to move the user to the actual test page, hosted at `https://browseraudit.com/test`. This intermediate step avoids surprising users by actively requiring their consent to begin the testing phase. Once the user clicks to start the tests, the main testing loop initiates.[2]

BrowserAudit is completely automated, and the user does not need to interact with the browser whilst it is being tested.

---

[1] For example, an exhaustive test of the same-origin policy would also need to demonstrate that, for any domains A and B, a page from domain A cannot access certain properties of a page from an incompatible domain B.

[2] Unless JavaScript is disabled, in which case we display a warning to the user. Automated tests cannot be run without JavaScript, and some security features need JavaScript in order to be exercised.

As the tests are running, the user can see a progress bar advancing, and four test counters being incremented, as shown in Figure 1. For the benefit of typical web users, test runs


Figure 1: The test summary box part-way through the execution of our tests.

are categorised using a simple Okay/Warning/Critical/Skipped traffic light indicator. Okay denotes passed tests, Warning and Critical denote failed tests, and Skipped denotes tests that are skipped because the feature being tested is not supported by the browser. Failures regarding SOP, cookies, and the `Referer` header, which we consider the most crucial security features, are reported as Critical; failures regarding CSP, CORS, HSTS and the `X-Frame-Options` header are reported as Warnings. This distinction is somewhat arbitrary, and will change as these features become more broadly supported and new ones are introduced.

After the test suite has finished running, the grey background of the summary box assumes the colour of the worst failed test, or green if all tests passed. This traffic light indicator provides a basic level of information about the current level of security offered by the browser.

More sophisticated users, such as security researchers or browser developers, need more information on the tests performed and on their outcomes. Clicking on the "Show/Hide Details" button displays a summary box that shows the various categories of tests (reflecting the security mechanisms that have been tested), and the number of failed tests for each of them, as shown in Figure 2.
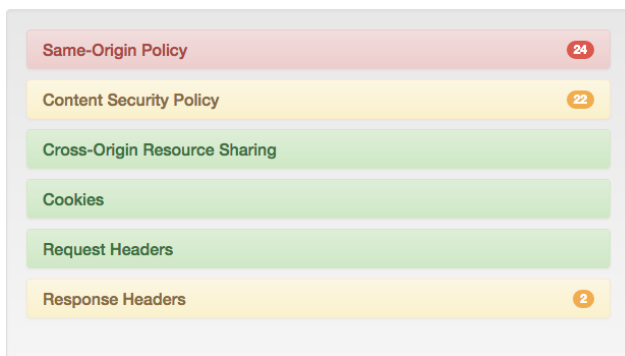
Figure 2: BrowserAudit summary box.

Each category can be expanded and collapsed to show a description of the corresponding security mechanism, and a list of sub-headers that in turn can be expanded to reveal individual tests for a specific feature, as illustrated in Figure 3. For each individual test we show a descriptive title that can

be clicked to show the client-side source code of the test itself. Our design uses the Bootstrap front-end framework [1], which
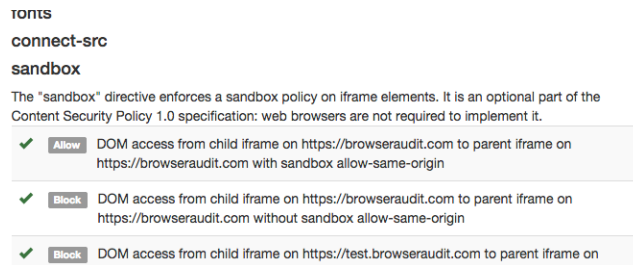
Figure 3: Some sub-categories of CSP tests, with expandable test titles and result indicators.

makes it easy to produce a layout that works consistently across browsers and devices.

2.2 Architecture

The client side and server side of BrowserAudit work together in order to run tests in the browser: the server side exercises browser security features, and the client side tests that these features are implemented as expected.

When multiple concurrent users access BrowserAudit, we need to avoid congestion on the server side, as testing each browser causes a bursty interaction with the BrowserAudit server in the form of hundreds of requests per user per minute. For this reason, we adopt a standard three-tier server architecture, consisting of a public-facing Nginx [11] web server, a Go [16] application server and a PostgreSQL [13] database backend. The Nginx server is running as a *reverse proxy* in front of the Go server, which is not publicly accessible. When the Nginx server receives HTTP(S) requests for static resources, such as our JavaScript tests, it responds by directly fetching the resource from the local `static/` directory. Dynamic requests are instead proxied to the Go server, and the responses are forwarded back to the client. Nginx also handles SSL termination, caching, gzip compression, URL rewriting, and keeps access and error logs. This architecture reduces the load on the Go server, which can focus on serving only dynamic requests that depend on the user's session, and limits security risks because the Go server can run as a non-privileged user.

Certificates. In order to ensure good coverage of various security features that involve the use of HSTS and cross-origin testing, BrowserAudit makes use of four domains: `browseraudit.com`, `test.browseraudit.com`, `browseraudit.org` and `test.browseraudit.org`. The server presents a single SSL certificate that is valid for all of these domains.

Sessions. We use sessions to keep track of intermediate test results and other test-related data for each user whilst their tests are in progress. Sessions are needed because in many of our security tests, it is the *server* that makes the decision as to whether or not the browser passed the test, not the test framework running in the browser. In these cases, the client must send an additional request asking the server what the test result was, so that it can be displayed to the user.

Caching. In our tests, there are many cases in which a request is first made to store a default result on the server, and then a second request *may* be sent to overwrite this result, depending on whether or not the browser correctly

implements a given security feature. If a user runs the tests twice in short succession, and this second result was cached and therefore did not reach our server, our application would report an incorrect test result. We ensure that this cannot happen by preventing HTTP responses from being cached.

2.3 Tests

A typical test of a security feature involves making multiple AJAX or image requests to the server and checking if the actual responses match the expected responses.

JavaScript and libraries. Our tests are written directly in JavaScript, using the jQuery library [9] for convenience. We deploy our tests using the Mocha framework for browser-based JavaScript unit testing [10], with some custom modifications to improve the output layout.

```
1  $.get("/del_httponly_cookie", function() {
2    expect($.cookie("httpOnlyCookie")).to.be.undefined;
3    $.get("/set_httponly_cookie", function() {
4      expect($.cookie("httpOnlyCookie")).to.be.undefined;
5      done();
6    });
7  });
```

Figure 4: The client side of a proof-of-concept `HttpOnly` cookie test.

Figure 4 shows a proof-of-concept test to check that the browser correctly implements `HttpOnly` cookies (see Section 3.4). Line 1 loads a page to clear any leftover cookies from previous test runs, line 2 checks that the cookie is not defined, line 3 loads a second page that sets the cookie, and line 4 checks that we are unable to read it via JavaScript. The call to `done()` on line 5 informs Mocha that the asynchronous test is complete. In order to make the source code of the tests easier to understand and maintain, we also leverage the Chai assertion library [7].

```
1  function ajaxSopTest(globalTestId, shouldBeBlocked,
        sourcePrefix, destPrefix) {
2    // omitted code: variable initialisation
3
4    var test_template = function(done) {
5      $.get("/sop/"+defaultResult+"/"+id,
6        function() {$("<iframe>", { src: iframeSrc })
7          .css("visibility", "hidden")
8          .appendTo("body").load(function() {
9            $.get("/sop/result/"+id,function(result) {
10             expect(result).to.equal("pass");
11             done();
12           });
13         });
14       });
15     };
16
17     // omitted code: save source code for display
18
19     browserAuditTest(globalTestId, test_template);
20  }
```

Figure 5: Code to generate SOP tests for AJAX calls.

Tests. In most cases, we automatically generate the JavaScript code for tests that have a similar structure but depend on different parameters. For example, in Figure 5 we show the most interesting parts of the `ajaxSopTest` function, which

generates Mocha code for testing AJAX calls with respect to the SOP. The choice of the right parameters for the resources to load (`defaultResults`, `iframeSrc`) are crucial to the correctness of each test instance. To favour modularity and coverage, we instantiate a separate Mocha test for each case to be tested, rather than bundling a large number of cases in the same test. To ensure maximum portability, we implement as much as possible on the client side using standard, browser-independent features.

Whenever possible, we write asynchronous tests using callback patterns rather than timeouts. We annotate the titles of tests whose results depend on timeouts with a small clock icon. We try to avoid using timeouts because, when a timeout expires, it is not possible to distinguish a true test failure from an anomalous delay in a browser event or network connection. Moreover, it is difficult to estimate appropriate timeout values for many events. For certain tests, however, we cannot avoid using timeouts. For example, to detect whether a CSP policy that denies the use of JavaScript but allows the loading of fonts in an iframe is enforced correctly, the BrowserAudit test framework needs to give time for the iframe to try to load the font, and then ask the server if the font was requested. We are not allowed to run JavaScript in the iframe to inspect the page and detect whether the font was loaded; likewise, we cannot ask the user for confirmation, because our tests must run without user interaction.

3. BROWSER SECURITY MECHANISMS

In this section, we describe the range of security mechanisms currently exercised by BrowserAudit. Each mechanism induces — sometimes implicitly — a security policy. Our emphasis is on testing representative instances of behaviours that should be allowed or blocked according to the corresponding security policy.

3.1 Same-Origin Policy

In the early days of the web, there was little incentive to control the resources that could be included in a web page: most web pages were static, and web developers were free to include resources (e.g., images) from any source in their web pages. As web sites became dynamic and interactive, thus allowing web developers to include user-supplied content in their pages, and requiring web browsers to execute scripts supplied by the web server, browser vendors became more security-aware: they recognised that permitting the execution of arbitrary code (e.g., JavaScript) from untrustworthy sources was potentially dangerous, and began to impose restrictions on the execution of scripts from "foreign" locations. In particular, "foreign" scripts were forbidden from accessing the Document Object Model (DOM) — the browser's internal hierarchical representation — of the web page in which the script was included. These are the foundations of the *same-origin policy* (SOP) [34], still implemented in contemporary web browsers: a script executing in the context of a web page is only permitted to access the DOM of another web page if the schemes, hostnames and port numbers in the URIs of the two pages — their *origins* — match.

There are mechanisms for relaxing the SOP so that information can be shared between DOMs with differing origins; the easiest method of doing so is to set the same `document.domain` property in each DOM, so that the web browser considers the DOMs to have the same origin.

BrowserAudit comprehensively exercises a web browser's

implementation of the SOP and the mechanisms for relaxing it to ensure that inter-DOM access is permitted when both DOMs are deemed to have the same origin, and is otherwise forbidden. Our DOM SOP tests have a common structure: scripts running on web pages loaded in nested iframes manipulate the DOM's `document.domain` property, and the script from one iframe attempts to access the DOM of the other iframe. Each test exercises a particular combination of the following parameters:

- The domain from which the web page loaded by the parent iframe is served (one of `browseraudit.{com/org}` or `test.browseraudit.{com/org}`);

- The domain from which the web page loaded by the child iframe is served (also selected from the list above, and potentially the same domain used by the parent iframe's web page);

- The value of `document.domain` to be set by a script running in the parent iframe;

- The value of `document.domain` to be set by a script running in the child iframe; and

- The direction in which the DOM access is attempted (parent iframe to child, or child iframe to parent).

The client-side test framework checks whether the web browser satisfies the SOP by selecting combinations of these parameters that should be allowed or blocked by the SOP and verifying that the correct behaviour is observed.

For example, Figure 6 shows a diagram for a test in which a parent iframe tries to access the DOM of its child iframe. The parent is loaded from `https://browseraudit.org` whereas the child is loaded from `https://test.browseraudit.org`. We expect this access to be blocked since we are not setting any `document.domain` values in this test, and the hostnames are not the same. To communicate test results from the server to the client whilst avoiding the restrictions imposed by the SOP itself, we use the established technique of loading images from specially-crafted addresses (i.e., `https://browseraudit.com/sop/[pass|fail]/TEST_ID`).

In general, if a script running in either iframe is able to access the DOM of the other, the script notifies the BrowserAudit server that access to the other iframe's DOM was granted; the test framework then queries the server for whether this notification was sent. If the notification was sent and DOM access was expected given the chosen test parameters, or if the notification was *not* sent and DOM access was *not* expected given the chosen test parameters, the test framework considers the browser to have passed that particular test; otherwise, the browser permitted insecure DOM access and is considered to have failed the test.

The SOP applies not only to DOM access, but also to cookies with differing paths and HTTP requests made to other domains via the XMLHttpRequest API; BrowserAudit also tests a browser's implementation of the SOP for all of these features, providing a total of 84 SOP tests, generated by four JavaScript templates.

3.2 Cross-Origin Resource Sharing

Cross-Origin Resource Sharing (CORS) [46] is a flexible standard for relaxing the SOP that selectively permits resources to be shared across origins; it is implemented in APIs capable of initiating cross-origin resource requests (e.g., XMLHttpRequest) in a range of modern web browsers. It allows a client to include a resource from a server with a different origin only if the resource request is explicitly authorised by the server. This is achieved via two additional HTTP headers: an `Origin` header is sent by the client as part of the request and specifies the origin of the resource attempting to use the cross-origin resource, and an `Access-Control-Allow-Origin` header is sent by the server as part of the response and specifies the origins from which this resource may be used, effectively ordering the client to uphold or relax the SOP for this resource request.

The majority of cross-origin requests made using CORS are "simple", defined in the CORS specification [46] as an HTTP request with one of `GET`, `POST` or `HEAD` as the request method and headers from a narrowly-defined whitelist (`Accept`, language-related headers and a small number of acceptable `Content-Type`s). Other requests are deemed "non-simple"; the CORS specification requires that the client precedes such requests with a "preflight" request that includes further detail so that the server can more accurately decide whether or not to allow the cross-origin request (although, in reality, some browsers misclassify simple and non-simple requests). In response to the preflight request, the server sends additional headers: `Access-Control-Allow-Methods`, a comma-delimited list of HTTP methods permitted to be used to access the resource; `Access-Control-Allow-Headers`, a comma-delimited list of headers that may be sent with the main CORS request; and `Access-Control-Expose-Headers`, a list of headers that should be exposed to the requester (e.g., a script accessing a resource using XMLHttpRequest). If the main CORS request violates either of the restrictions imposed by the `Access-Control-Allow` headers, the main request is considered a violation of the SOP and is aborted.

BrowserAudit exercises the browser's implementation of CORS by sending a series of cross-origin XMLHttpRequest requests from the browser and verifying that the client exhibits CORS-compliant behaviour when the BrowserAudit server sends a response containing a range of CORS HTTP headers. The testing methodology is similar to that for the SOP, described in Section 3.1: the client attempts to retrieve a file from the BrowserAudit server, and sends a notification to the BrowserAudit server if this retrieval was successful. The BrowserAudit test framework then queries the server for whether the notification was sent. If the notification was sent for CORS-compliant requests and *not* sent for CORS-violating requests, the browser is deemed to correctly implement the CORS standard; if a notification was sent for CORS-violating requests, or if one was *not* sent for CORS-compliant requests, the browser is considered to lack full compliance.

We currently test 54 different CORS scenarios, automatically generated by four JavaScript test templates.

3.3 Content Security Policy

The *Content Security Policy* (CSP) standard[3] [45] enables much finer-grained control over the loading of arbitrary re-

---

[3] We concern ourselves only with version 1.0 of the Content Security Policy standard, as its successor (version 1.1) is still in Working Draft status at the time of writing; however, the two versions are similar, and the latter can be viewed as an extension of the former.
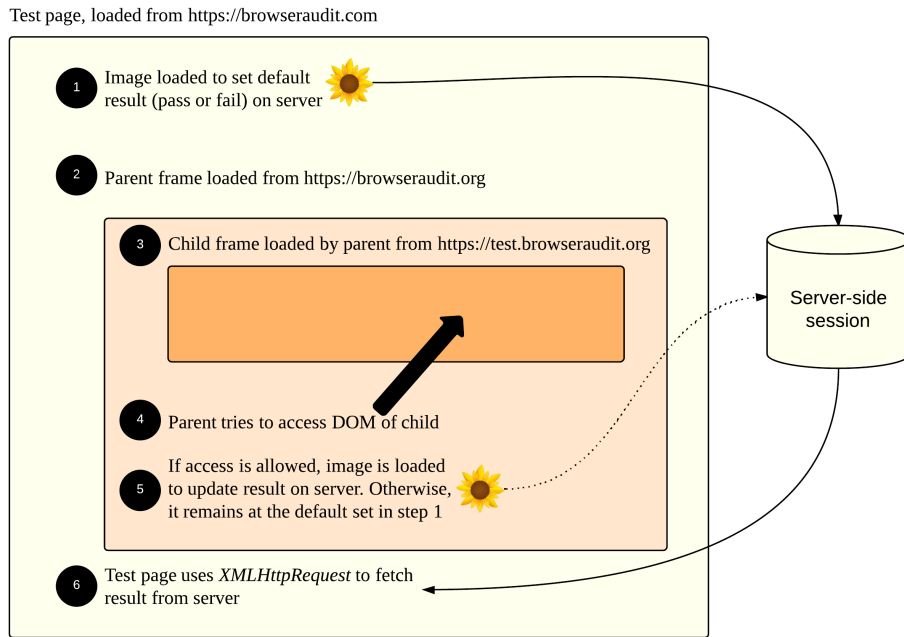
Test page, loaded from https://browseraudit.com

1 Image loaded to set default result (pass or fail) on server

2 Parent frame loaded from https://browseraudit.org

3 Child frame loaded by parent from https://test.browseraudit.org

4 Parent tries to access DOM of child

5 If access is allowed, image is loaded to update result on server. Otherwise, it remains at the default set in step 1

6 Test page uses *XMLHttpRequest* to fetch result from server

Server-side session

Figure 6: An example of an SOP test in which the parent frame tries to access the DOM of its child.

sources on a web page than the SOP and CORS. As with CORS, a content security policy is delivered via an HTTP header (or via a `<meta>` element in the HTML header); the CSP specification states that the `Content-Security-Policy` header should be used for this purpose.

The header allows servers to declare to CSP-compliant clients the permitted origins of a range of resources: images, stylesheets, scripts, web fonts, embedded objects and other types of resource may all be controlled by a single policy. *Directives* may be used to restrict the origins of these different types of resource independently of each other, and a "default" directive may be used to restrict the origins of all resources that are not explicitly controlled elsewhere in the policy. For example, a server at `example.com` serving web pages to CSP-compliant browsers could restrict the loading of images to those hosted on the same server and the loading of embedded objects (such as Java applets) to those hosted on a trusted server at `applets.example.com` (and thus forbid embedded objects and images from being loaded from other origins) by specifying the following value for the `Content-Security-Policy` header:

`image-src 'self'; object-src http://applets.example.com`

When served alongside a web page to a CSP-compliant web browser, such policies can preempt many common web attacks; e.g., using the `script-src` directive to control the permissible origins of scripts mitigates the effects of CSRF, clickjacking and framebusting (since they rely primarily on successful JavaScript injection), and using the `style-src` directive to control the permissible origins of stylesheets defeats CSS-based attacks. Note that one cannot specify which specific resources may be loaded from these other origins: permitting a particular Java applet to be loaded from `applets.example.com` also permits *any other* embeddable

object to be loaded from `applets.example.com`, so whitelisted origins should be trustworthy (particularly those granting the power to execute arbitrary code, such as `script-src`).

The CSP standard also includes a mechanism for reporting violations of a given policy via a special `report-uri` directive; this directive defines a URL to which a *violation report* should be sent.

BrowserAudit exercises a browser's CSP implementation by performing a battery of tests on each directive defined in the CSP specification, as well as the violation-reporting capabilities of the `report-uri` directive. Similarly to the SOP tests (described in Section 3.1), each CSP test attempts to load a resource inside an iframe using a particular combination of the following parameters:

- The domain from which the web page loaded by the iframe is served (one of `browseraudit.com` or `test.browseraudit.com`);

- The domain from which the desired resource is requested (also selected from the list above, and potentially the same domain used by the iframe's web page); and

- The CSP imposed on the iframe by the BrowserAudit server via the `Content-Security-Policy` header.

We run 226 CSP tests, generated by three JavaScript templates, that in turn load approximately 280 iframes representing particular behaviours to be tested. In each test, the browser is expected to either allow or block access to the given resource, and the act of requesting the resource from the BrowserAudit server allows it to track violations of the given policy. On the client side, the BrowserAudit test framework queries the server after the iframe has loaded to find whether the browser accessed the resource and therefore determine whether the browser exhibited the behaviour

expected of a CSP-compliant browser: allowing a request permitted by the given policy or blocking a request restricted by the policy is regarded as a correct implementation of the CSP standard and thus a test success, whilst an attempt to access the resource when given a restrictive policy or a failure to request the resource when given a permissive policy is regarded as an erroneous implementation of the standard and thus a test failure.

```
1 $("<iframe>", { src: "/csp/serve/206/param-html?policy='
       sandbox allow-same-origin allow-scripts'&defaultResult=
       pass" })
2   .css("visibility", "hidden").appendTo("body")
3   .load(function() {
4     $.get("/csp/result/206", function(result) {
5       expect(result).to.equal("pass");
6       done();
7     });
8   });
```

Figure 7: A CSP test exercising the browser's implementation of the sandbox directive.

```
1 <html><body>
2   <iframe src="/csp/serve/206/param-htmlb?sessid=
          sessionCookie"></iframe>
3 </body></html>
```

Figure 8: The HTML for the outer iframe loaded by the test script shown in Figure 7.

Figure 7 shows the client-side code for a CSP test. The code runs on the main BrowserAudit page and loads an outer iframe from browseraudit.com with the CSP header sandbox allow-same-origin allow-scripts. This outer iframe is very simple (Figure 8), and its role is simply to load an inner iframe from browseraudit.com that is subject to the given policy: scripts can run, and have same-origin permissions. The inner frame, whose code is shown in Figure 9, tries to perform an XMLHttpRequest to test.browseraudit.com, which should be blocked. Note that since we cannot rely on user credentials to be sent with synchronous XMLHttpRequests, we pass the session cookie (abstracted for readability in Figure 9 as sessionCookie) as a parameter of the request. All of this information is also visible to the BrowserAudit user by clicking on the corresponding test title in the user interface.

```
1 <html><body>
2   <script>
3     var xhr = new XMLHttpRequest();
4     xhr.open("GET", "https://test.browseraudit.com/csp/serve
          /206/oktext?sessid=sessionCookie&corsOrigin=
          browseraudit.com&corsMethod=GET", false);
5     xhr.send(null);
6     if (xhr.status == 200) {
7       var img = document.createElement("img")
8       img.setAttribute("src", "/csp/fail/206/png");
9       document.body.appendChild(img);
10    }
11  </script>
12 </body></html>
```

Figure 9: The HTML for the inner iframe corresponding to the outer iframe shown in Figure 8.

3.4 Cookies

In our SOP tests (Section 3.1) we explore the security implications of setting the cookie scope through the Domain and Path attributes. There are two other important aspects of cookie security: the HttpOnly and Secure attributes. We test the browser's treatment of these attributes, expecting the behaviour defined in RFC 6265 [17].

The HttpOnly attribute of a cookie instructs the browser to reveal that cookie only through an HTTP request; i.e., it should not be made available to client-side scripts. The benefit of this is that, even if an XSS vulnerability is exploited, the cookie cannot be stolen. HttpOnly cookies are supported by all major browsers, with the notable exception of Android 2.3's stock browser. BrowserAudit includes tests that check that an HttpOnly cookie sent from the server cannot then be accessed by JavaScript, and that HttpOnly cookies cannot be created by JavaScript.

When a cookie has the Secure attribute set, a compliant browser will include the cookie in an HTTP request only if the request is transmitted over a secure channel (i.e., in an HTTPS request). This keeps the cookie confidential: an attacker would not be able to read it even if he were able to intercept the connection between the victim and the destination server. The Secure attribute is supported by all major browsers. BrowserAudit includes tests checking the browser's treatment of the Secure attribute both when the cookies are set by the server and set by JavaScript.

3.5 Referer Header

The Referer header should not be included in a non-secure request if the referring page was served via a secure protocol; this behaviour is defined in RFC 2616 [31]. This requirement exists because the referrer might disclose an otherwise private information source. In BrowserAudit, we test this behaviour by loading a web page over HTTPS containing an image loaded over HTTP and checking that the Referer header was not sent to the server with the request for the image.

3.6 Response Headers

3.6.1 X-Frame-Options

X-Frame-Options, defined in RFC 7034 [42], is a server-side technique that can be used to prevent clickjacking attacks. X-Frame-Options is a response header that specifies whether or not the document being served is allowed to be rendered in a frame; more specifically, the header specifies the origin (scheme, hostname and port number) that is allowed to render the document in a frame. BrowserAudit tests for correct treatment of the DENY, SAMEORIGIN and ALLOW-FROM directives. The tests try to load iframes served with different headers; each iframe that loads reports its success to the server, which assesses whether the browser behaved as expected. Our tests currently only cover the <iframe> element, although the header also applies to <frame>, <object>, <applet> and <embed> elements.

X-Frame-Options is supported in all modern browsers, although the implementations across browsers differ. Some browsers behave differently when dealing with nested frames, so we do not test these cases as there is no defined correct behaviour. Note also that not all browsers support the ALLOW-FROM directive.

3.6.2 Strict-Transport-Security

HTTP Strict Transport Security (HSTS) is a security mechanism that allows a server to instruct browsers only to communicate with it over a secure (HTTPS) connection for the given domain. It exists primarily to defend against man-in-the-middle attacks in which an attacker is able to intercept his victim's network connection [37]. The server sends this instruction via the `Strict-Transport-Security` header, as defined in RFC 6797 [35].

When HSTS is enabled on a domain, a compliant browser must rewrite any plain HTTP requests to that domain to use HTTPS. This includes both URLs entered into the navigation bar by the user, and resources included on a web page. The `Strict-Transport-Security` header should only be sent in an HTTPS response. If the browser receives the header in a response sent over plain HTTP, it should be ignored.

In BrowserAudit, we test the basic behaviour of HSTS and its `includeSubDomains` directive. We also ensure that the header is ignored when transferred via an insecure protocol, and that the HSTS state correctly expires based on the `max-age` value set in the header. All of these tests work by testing whether a request for an image at `http://browseraudit.com/set_protocol` is rewritten to use HTTPS or not.

Almost all current browsers support HSTS, with the notable exception of Internet Explorer 11 (the latest available version at the time of writing).

4. EVALUATION

4.1 Performance

A primary concern of BrowserAudit is scalability, given that a single invocation of the full test suite invokes approximately 1,500 requests and transfers around 3MB of data between the client and server. The server must handle all of these requests quickly (ideally in under 300ms), given the large number of tests in the BrowserAudit test suite and the reliance of some of the tests on timeouts (see Section 2.3).

The BrowserAudit web and database servers are currently hosted on a single virtualised server with two CPU cores and 2GB of memory, running Ubuntu 14.04. We evaluated BrowserAudit's server-side performance by running the BrowserAudit test suite in 15 web browsers repeatedly and concurrently for 15 minutes. Over this period, the BrowserAudit server handled around 225,000 requests and served a total of 450MB of data. The 1- and 5-minute load averages on the BrowserAudit server are shown in Figure 10; the peak load averages over the 15-minute duration of the performance test are 1.2 and 0.7 respectively, where a load average of 1 indicates that a single CPU core is operating at capacity. Based on these performance figures, we estimate that a single BrowserAudit application server using this configuration could comfortably support up to 25 concurrent test suite executions.

As described in Section 2.2, our design is ready to be scaled up as the BrowserAudit user base grows. Nginx can be configured as a load balancer, passing requests to one of many application servers. Deploying Go application server instances is trivial thanks to Go's ability to compile a program to a single statically-linked binary, so there is no dependency chain. In order to maintain session persistence, Nginx's `ip_hash` directive can be used to ensure that all requests from the same IP address reach the same application server, maintaining the integrity of a single suite execution.
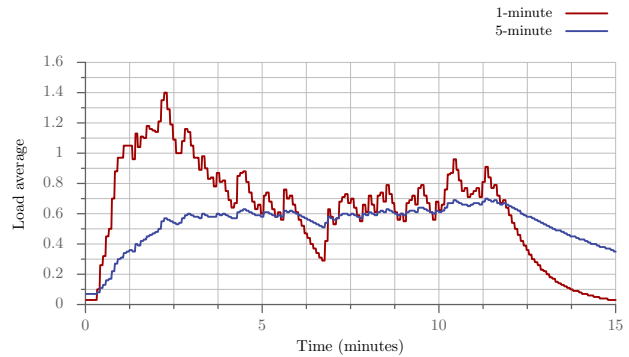

Figure 10: The 1- and 5-minute load averages on the BrowserAudit server during the performance evaluation.

Most client-side tests contain components that are loaded synchronously inside dynamically-created iframes, which become redundant as soon as the test result is reported in the browser; over time, the DOM of the main BrowserAudit window would therefore amass an overwhelming number of iframes, slowing down the execution of tests as the browser struggles to create and append additional iframes. We avoid this problem by dynamically removing any iframes appended to the DOM during each test's tear-down phase (via Mocha's `afterEach()` routine). We ran 15 repetitions of 10 concurrent executions of the whole test suite on a 64-bit Windows 7 machine with a 6-core Intel i7 4930K CPU and 64GB of memory, and Chromium 40.0.2205.0. Under these conditions, the average execution time for the test suite is just over a minute. By contrast, a single execution in Safari 8.0 on an iPhone 5 with iOS 8.1 takes on average 1.35 minutes, skipping 24 tests. The execution time varies broadly across browsers and platforms, but we consider this an acceptable cost for performing an in-depth browser security scan.

4.2 Correctness

Verifying the correctness of our tests is challenging, as they need to convey in a final pass or fail result a whole security-sensitive behaviour: a test containing a small bug could still pass, which is generally the expected result for browsers correctly implementing a given security mechanism.

Of course, no web browsers contain intentional security flaws that would allow us to verify the correctness of tests. Modifying the source code of existing open-source browsers to break their security features in order to ensure that tests fail when expected is possible but challenging given the complexity of modern web browser code bases.

However, it is a matter of public record that some web browsers either do not implement some of the security mechanisms tested by BrowserAudit, or only implement subsets of those security mechanisms. We leverage the results of browser-profiling projects such as *Browserscope* [3] and *Can I Use. . .* [6] to broadly identify the security features implemented by each web browser, and for those features we manually verify that the BrowserAudit test suite results are accurate.

Using *BrowserStack* [5], a web-based browser testing service, we have evaluated BrowserAudit in a range of browsers on a number of different operating systems, across both

desktop and mobile platforms. The full BrowserAudit test suite runs reliably in Safari 6, Firefox 13 and Chrome 25 or more recent versions, automatically skipping tests where a feature is not supported. BrowserAudit also runs correctly on Internet Explorer 11, but due to problems relating to Mocha and IE's limited call stack, it cannot execute the whole test suite. In older versions of these browsers, it is instead possible to run a subset of the test suite.

4.3 Test coverage

We noted in Section 2 that full coverage for browser security feature tests is unattainable. Here we discuss a number of security features not covered by BrowserAudit, but that we believe can be added to our framework.

We imply in Section 3 that there is no single same-origin policy but rather a collection of related security mechanisms. We currently test the same-origin policy for DOM access, XMLHttpRequest and cookies. This could be expanded to test the same-origin policies for Flash, Java, Silverlight, and HTML5 web storage.

The postMessage API is used by many developers to communicate across origins [19]. Since the API allows the sender of a message to specify the origins of the recipients that may receive the message, there are lots of origin-related tests that we could write for this feature in BrowserAudit.

Another security feature that could be tested is the X-Content-Type-Options response header first introduced in Internet Explorer 8 [40]. It is now also supported by Chromium and Safari; the Firefox team is still debating its implementation [26]. It is designed to prevent browser-sniffing attacks where a resource (e.g., a HTML document) is sent with an inappropriate MIME type (e.g., text/plain) but is nonetheless erroneously rendered by the browser as if the correct MIME type had been sent [18].

In Sections 3.1–3.4 we discussed how to extend coverage of features for which we already have some tests. Summarising, the main limitations are that: in many tests involving origin mismatches, we only test origins that differ by hostname rather than by scheme or port number; we do not test CSP directives where a resource is loaded from a URL that redirects; we do not test that cookies cannot be set for top-level domains that include a country code, such as co.uk (whereas, for example, they should be settable for example.uk). We also do not test the Report-Only header defined by the CSP standard, but this is not due to a limitation of the BrowserAudit framework and a suitable test could be added to the test suite.

Finally, cryptographic APIs such as the W3C WebCrypto API and the OpenSSL library are important aspects of browser security, but cryptographic testing is beyond the scope of BrowserAudit and better left to dedicated projects such as *How's My SSL?* [8].

4.4 Uncovering Security Bugs

BrowserAudit's test suite has uncovered two previously-unknown bugs in Firefox's implementation of the CSP standard; these bugs are present in all versions of Firefox that implement the CSP standard up to version 32.0.3. The first bug [24] allows the loading of same-origin stylesheets with the policy

```
default-src 'none'; style-src 'unsafe-inline';
```

similarly, the second bug [25] allows the loading of same-origin Worker and SharedWorker objects in scripts with the policy

```
default-src 'none'; script-src 'unsafe-inline'.
```

In both cases, the 'unsafe-inline' declaration in the policy states that only inline stylesheets and scripts must be permitted: external resources, even those from the same origin, must be blocked. We reported both of these bugs to Mozilla during the version 29 release cycle, and they were fixed in version 33 of Firefox.

Firefox does not currently implement the sandbox CSP directive; this optional feature of the CSP 1.0 specification directs browsers to relax the given security controls on iframes embedded in the page, as if they had been supplied in the sandbox attribute of each <iframe> element. The sandbox attribute is in fact a feature of the HTML5 specification [34] and states that an iframe containing a sandbox attribute should have *all* security controls enabled unless specifically disabled by values inside the sandbox attribute. Development work on the implementation of this directive in Firefox is currently underway [27]. However, the current implementation does not correctly handle the case where an empty value is given for the sandbox CSP directive; the CSP 1.0 specification implies that the browser should apply a sandbox attribute with an empty value (and thus enforce a highly-restrictive sandboxing policy — a view also taken by developers of other browsers, such as Chromium), but Firefox's implementation does not apply a sandbox attribute at all in this scenario (thus failing to enforce *any* sandboxing policy). This flaw was uncovered by the current set of CSP tests in BrowserAudit, and we are in discussions with Firefox developers to address it before their sandbox implementation lands in a stable version of the browser.

5. RELATED WORK

In this section we discuss some related work on browser security, which influenced the design of our tests, and review some web applications that perform security-relevant tests, which served as a source of inspiration for BrowserAudit.

5.1 Browser Security

The authoritative sources of information on upcoming browser security mechanisms are of course the W3C RFCs and Drafts such as [34, 45, 21, 46, 35]. Most security measures are the result of a lot of practical experimentation and academic research that led to proposals that gradually gained adoption and became more robust through security reviews and public scrutiny. Paradigmatic examples are the early contributions of Barth, Jackson et al. to postMessage, the Origin header and HTTPS [19, 37, 20].

The standards themselves provide a lot of detail about the intended security behaviour, but additional research is needed to interpret the consequences for deployed web applications. For example, De Ryck et al. perform a security analysis of some of the upcoming standards in [28], finding them to be be of high quality but also highlighting potential security risks. Singh et al. [44] discover potentially dangerous incoherencies amongst different browser access control policies.

A broad, in-depth analysis of browser security can be found in Zalewski's *Browser Security Handbook* [47] and the companion book *The Tangled Web* [48]; they gather a wealth of information on browser security features, their shortcomings and the peculiar differences in browser support.

5.2 Web Sites

Panopticlick [12] is an experiment to investigate how unique — and therefore trackable — modern web browsers are, by fingerprinting their version and configuration information. Some of this information can be gleaned directly from browser requests, whereas other information is made available by the presence of JavaScript and browser plugins. Visitors click a "Test Me" button and are then provided with their browser's uniqueness score and a breakdown of the measurements used to obtain the result. These data are then anonymously stored in the project database to make future uniqueness scores more accurate, and to allow for analysis of the data, as discussed in [30]. Although focussed on privacy rather than security, *Panopticlick* was the main inspiration for BrowserAudit.

BrowserSpy [4] is another web site that reports how much information can be retrieved from a browser by visiting a test page. Its focus is on privacy, yet some of its tests are security-related, although not presented as such; for example, one test checks that JavaScript cannot read `HttpOnly` cookies. Each of *BrowserSpy*'s 75 current tests has to be run individually, since the output is rather verbose, and the output does not show implementation details that could be useful for a technical audience. In contrast, our 400+ tests run automatically, and advanced users can view the client-side code driving each individual test.

How's My SSL? [8] is a recent project that advises the user on the security of their TLS client (web browsers act as TLS clients when engaged in HTTPS communication). It works by running a TLS server that has been modified so that the client-server handshake is exposed to the web application, allowing it to inspect the cipher suites that the client supports and perform a security assessment. The results are reported clearly, with "Learn More" links for more technical background which also inspired our design. The test results can be accessed via a JSON API, and could be potentially integrated into BrowserAudit to complement our tests. *Qualys SSL Labs* [14] also offers browser-based tests for SSL clients that display a concise report of their TLS capabilities, intended for the expert user. In BrowserAudit we instead strived to produce reports that can be interpreted by users at different levels of technical competence.

The *Can I Use...* test suite [15] gathers browser compatibility data for a wide variety of browser features such as support for HTML5 and CSS3. Some of these tests are automatic and others require visual confirmation or interaction from the user. A few tests check for support for security features; for example, one (interactive) test detects support for the CSP. In contrast, BrowserAudit runs 226 automated tests to assess the security of the CSP implementation.

The *Browser DOM access checker* [2] is a web page also included in the Chromium browser source code that uses JavaScript to test the enforcement of some domain-related security policies such as cross-domain DOM access, JavaScript cookie access, `XMLHttpRequest` calls, and event and transition handling; for example, it runs hundreds of tests to ensure that read or write attempts to the visible properties of the `document` object are blocked cross-domain. In contrast, we are satisfied with testing cross-domain access for one representative property of the `document` object: if such access is blocked, we conclude that the policy is effective. We could programmatically extend our tests to try accessing all properties, but that goes beyond the scope of BrowserAudit:

DOM-based cross-domain access is only one of the hundreds of qualitatively different behaviours that we consider.

Finally, *Browserscope* [3] is a community-driven project for profiling web browsers; it detects the browser version and runs tests that cover a broad range of features such as network performance, CSS support, and JavaScript optimisations. Test results are aggregated and made publicly available, making it easy for web developers to keep track of functionality across all browsers that have been tested.

Currently, *Browserscope* also includes 17 tests which automatically check whether the browser supports a number of standard features relevant to security and displays a list of which tests passed or failed. In contrast, BrowserAudit is engineered to run hundreds of tests that ascertain whether security features are implemented correctly, and provides an interface that allows different types of users to access detailed descriptions of each test case, including client-side test code.

6. CONCLUSIONS

We introduced BrowserAudit, a web application to test the implementation of browser security features. It complements the unit testing used by browser vendors to debug their implementations by checking that deployed browsers effectively deliver the security behaviours entailed by the specifications of browser security mechanisms.

All of our tests run automatically without interaction from the user, and provide detailed information for each test category, including the source code of each individual test. This makes BrowserAudit useful for a broad audience, from the casual user to the web developer and the security researcher. No other publicly-accessible web application tests such a breadth of browser security mechanisms as ours, either established or experimental.

In Section 4.3 we highlighted aspects of browser security mechanisms that are currently not covered by our tests. BrowserAudit is designed to be modular and extensible; adding variants of existing tests with different combinations of parameters, or new client-side-only tests (e.g., to test different features of the SOP) is straightforward. We are currently investigating the more challenging problem of allowing similar extensibility of the server-side components of tests.

BrowserAudit is an open-source project [32], and we hope that the web security community will help us extend it with even more test cases.

7. ACKNOWLEDGMENTS

8. REFERENCES

[1] Bootstrap. http://getbootstrap.com/.
[2] Browser DOM access checker. http://lcamtuf.coredump.cx/dom_checker/.
[3] Browserscope. http://www.browserscope.org/.
[4] BrowserSpy. http://browserspy.dk/.
[5] BrowserStack. http://www.browserstack.com/.
[6] Can I Use.... http://caniuse.com/.
[7] Chai. http://chaijs.com/.

[8] How's My SSL? https://www.howsmyssl.com/.

[9] jQuery. http://jquery.com/.

[10] Mocha. http://mochajs.org/.

[11] Nginx. http://nginx.org/.

[12] Panopticlick. https://panopticlick.eff.org/.

[13] PostgreSQL. http://www.postgresql.org/.

[14] Qualys SSL Labs. https://www.ssllabs.com/.

[15] The Can I Use... test suite. http://tests.caniuse.com/.

[16] The Go Programming Language. https://golang.org/.

[17] A. Barth. HTTP State Management Mechanism. RFC 6265 (Proposed Standard), Apr. 2011.

[18] A. Barth, J. Caballero, and D. Song. Secure content sniffing for web browsers, or how to stop papers from reviewing themselves. In *Proceedings of S&P 2009*, pages 360–371, 2009.

[19] A. Barth, C. Jackson, and J. Mitchell. Securing Frame Communication in Browsers. In *Proceedings of USENIX Security 2008*, pages 17–30, 2008.

[20] A. Barth, C. Jackson, and J. C. Mitchell. Robust Defenses for Cross-site Request Forgery. In *Proceedings of CCS'08*, pages 75–88, 2008.

[21] A. Barth and M. West. Content Security Policy 1.1, June 2013. W3C Working Draft WD-CSP11-20130604.

[22] K. Bhargavan, A. Delignat-Lavaud, and S. Maffeis. Language-Based Defenses Against Untrusted Browser Origins. In *Proceedings of USENIX Security 2013*, pages 653–670, 2013.

[23] E. Budianto, Y. Jia, X. Dong, P. Saxena, and Z. Liang. You Can't Be Me: Enabling Trusted Paths and User Sub-origins in Web Browsers. In A. Stavrou, H. Bos, and G. Portokalidis, editors, *Proceedings of RAID 2014*, volume 8688 of *Lecture Notes in Computer Science*, pages 150–171. Springer, 2014.

[24] Bugzilla. Bug 1007205 — CSP allows local CSS @import with only 'unsafe-inline' set. https://bugzilla.mozilla.org/show_bug.cgi?id=1007205.

[25] Bugzilla. Bug 1007634 — CSP allows local Worker construction with only 'unsafe-inline' set. https://bugzilla.mozilla.org/show_bug.cgi?id=1007634.

[26] Bugzilla. Bug 471020 — Add X-Content-Type-Options: nosniff support to Firefox. https://bugzilla.mozilla.org/show_bug.cgi?id=471020.

[27] Bugzilla. Bug 671389 — Implement CSP sandbox directive. https://bugzilla.mozilla.org/show_bug.cgi?id=671389.

[28] P. De Ryck, L. Desmet, P. Philippaerts, and F. Piessens. A security analysis of next generation web standards. Technical report, ENISA, July 2011.

[29] X. Dong, Z. Chen, H. Siadati, S. Tople, P. Saxena, and Z. Liang. Protecting sensitive web content from client-side vulnerabilities with CRYPTONS. In *Proceedings of CCS'13*, pages 1311–1324, 2013.

[30] P. Eckersley. How unique is your web browser? In *Proceedings of PETS'10*, pages 1–18, 2010.

[31] 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), June 1999.

[32] GitHub. BrowserAudit project. https://github.com/browseraudit/.

[33] M. Heiderich, M. Niemietz, F. Schuster, T. Holz, and J. Schwenk. Scriptless Attacks: Stealing the Pie Without Touching the Sill. In *Proceedings of CCS'12*, pages 760–771, 2012.

[34] I. Hickson and D. Hyatt. HTML5: A vocabulary and associated APIs for HTML and XHTML. W3C Candidate Recommendation CR-HTML5-20140429, Apr. 2014.

[35] J. Hodges, C. Jackson, and A. Barth. HTTP Strict Transport Security (HSTS). RFC 6797 (Proposed Standard), Nov. 2012.

[36] L.-S. Huang, A. Moshchuk, H. J. Wang, S. Schechter, and C. Jackson. Clickjacking: Attacks and Defenses. In *Proceedings of USENIX Security 2012*, pages 22–22, 2012.

[37] C. Jackson and A. Barth. Forcehttps: Protecting High-security Web Sites from Network Attacks. In *Proceedings of WWW'08*, pages 525–534, 2008.

[38] E. Kirda. Cross Site Scripting Attacks. In *Encyclopedia of Cryptography and Security*, pages 275–277. 2011.

[39] S. Maffeis, J. C. Mitchell, and A. Taly. Object Capabilities and Isolation of Untrusted Web Applications. In *Proceedings of S&P 2010*, pages 125–140, 2010.

[40] MSDN Blogs. IE8 Security Part VI: Beta 2 Update. http://blogs.msdn.com/b/ie/archive/2008/09/02/ie8-security-part-vi-beta-2-update.aspx.

[41] K. Patil, X. Dong, X. Li, Z. Liang, and X. Jiang. Towards Fine-Grained Access Control in JavaScript Contexts. In *Proceedings of ICDCS'11*, pages 720–729, 2011.

[42] D. Ross and T. Gondrom. HTTP Header Field X-Frame-Options. RFC 7034 (Informational), Oct. 2013.

[43] G. Rydstedt, E. Bursztein, D. Boneh, and C. Jackson. Busting Framebusting: a Study of Clickjacking Vulnerabilities at Popular Sites. In *Proceedings of W2SP 2010*, 2010.

[44] K. Singh, A. Moshchuk, H. J. Wang, and W. Lee. On the Incoherencies in Web Browser Access Control Policies. In *Proceedings of S&P 2010*, pages 463–478, 2010.

[45] B. Sterne and A. Barth. Content Security Policy 1.0. Nov. 2012. W3C Candidate Recommendation CR-CSP-20121115.

[46] A. Van Kesteren. Cross-origin Resource Sharing. W3C Recommendation REC-cors-20140116, Jan. 2014.

[47] M. Zalewski. Browser Security Handbook, 2010.

[48] M. Zalewski. *The Tangled Web: A Guide to Securing Modern Web Applications.* No Starch Press, 2012.