

WebPol: Fine-grained Information Flow Policies for Web Browsers

Abhishek Bichhawat¹, Vineet Rajani², Jinank Jain³, Deepak Garg², and Christian Hammer⁴

¹ Saarland University, Germany

² MPI-SWS, Germany

³ ETH Zürich, Switzerland

⁴ University of Potsdam, Germany

Abstract. In the standard web browser programming model, third-party scripts included in an application execute with the same privilege as the application’s own code. This leaves the application’s confidential data vulnerable to theft and leakage by malicious code and inadvertent bugs in the third-party scripts. Security mechanisms in modern browsers (the same-origin policy, cross-origin resource sharing and content security policies) are too coarse to suit this programming model. All these mechanisms (and their extensions) describe whether or not a script can access certain data, whereas the meaningful requirement is to allow untrusted scripts access to confidential data that they need and to prevent the scripts from leaking data on the side. Motivated by this gap, we propose WebPol, a policy mechanism that allows a website developer to include fine-grained policies on confidential application data in the familiar syntax of the JavaScript programming language. The policies can be associated with any webpage element, and specify *what* aspects of the element can be accessed *by which third-party domains*. A script can access data that the policy allows it to, but it cannot pass the data (or data derived from it) to other scripts or remote hosts in contravention of the policy. To specify the policies, we expose a small set of new native APIs in JavaScript. Our policies can be enforced using any of the numerous existing proposals for information flow tracking in web browsers. We have integrated our policies into one such proposal that we use to evaluate performance overheads and to test our examples.

1 Introduction

Webpages today rely on third-party JavaScript to provide useful libraries, page analytics, advertisements and many other features. JavaScript works on a *mashup* model, wherein the hosting page and included scripts share the page’s state (called the DOM). Consequently, by design, all included third-party scripts run with the same access privileges as the hosting page. While some third-party scripts are developed by large, well-known, trustworthy vendors, many other scripts are developed by small, domain-specific vendors whose commercial motives do not always align with those of the webpage providers and users. This

leaves sensitive information such as passwords, credit card numbers, email addresses, click histories, cookies and location information vulnerable to inadvertent bugs and deliberate exfiltration by third-party scripts. In many cases, developers are fully aware that a third-party script accesses sensitive data to provide useful functionality, but they are unaware that the script also leaks that data on the side. In fact, this is a widespread problem [18].

Existing web security standards and web browsers address this problem unsatisfactorily, favoring functionality over privacy. The same-origin policy (SOP) [5] implemented in all major browsers restricts a webpage and third-party scripts included in it to communicating with web servers from the including webpage’s domain only. However, broad exceptions are allowed. For instance, there is no restriction on request parameters in urls that fetch images and, unsurprisingly, third-party scripts leak information by encoding it in image urls. The candidate web standard Content Security Policy (CSP) [31], also implemented in most browsers, allows a page to white list scripts that may be included, but places no restriction on scripts that have been included, thus not helping with the problem above. Other mechanisms (including a provision in the SOP) restrict scripts loaded in a different third-party window or frame from accessing the resources of a page but do not restrict third-party scripts included in the page itself.

The academic community has recently proposed solutions based on information flow control (IFC) [12,30,16,7,26,9,28,18], also known as mandatory access control. Their *ideal* goal is to allow third-party scripts access to necessary sensitive data, but restrict *where* scripts can send the data—and data derived from that data—in accordance with a policy. While this would balance functionality and privacy perfectly, all existing IFC-based solutions for web browsers fall short of this ideal goal. Many proposals, including several taint-based solutions [16,7,26,9], focus on the IFC mechanism, but currently lack adequate support for specifying policies conveniently. Flowfox [30] provides a rich policy framework but all websites are subject to the same policy, and the underlying IFC technique, secure multi-execution [13], does not handle shared state soundly. COWL [28] uses coarse-grained isolation, allowing scripts’ access to either remote domains or the shared state, but not both. This requires significant code changes when both are needed simultaneously (see Section 7 for more details).

The contribution of our work is WebPol, a policy framework that allows a webpage developer to release data selectively to third-party scripts (to obtain useful functionality), yet control what the scripts can do with the data. WebPol integrates with any taint-based IFC solution to overcome the shortcomings listed above. WebPol policies label sensitive content (page elements and user-generated events) at source, and selectively declassify them by specifying where (to which domains) the content and its derivatives can flow. Host page developers specify WebPol policies in JavaScript, a language already familiar to them.

Under the hood, any taint-based IFC solution can be used to track data flows and to enforce WebPol policies. As a demonstrative prototype, we have integrated WebPol with our previous taint-based IFC framework for WebKit [7,26], the engine that powers Apple’s Safari and other browsers. We demonstrate the

expressiveness of WebPol policies through examples and by applying WebPol to two real websites. Through measurements, we demonstrate that WebPol policies impose low-to-moderate overhead, which makes WebPol usable today. A full version of the paper [8] contains more details, including a small lab study through which we test that WebPol can be effectively used by programmers familiar with HTML and JavaScript.

2 Overview

This section provides an overview of information flow control (IFC) in the context of web browsers and lists important considerations in the design of WebPol. IFC is a broad term for techniques that control the flow of sensitive information in accordance with pre-defined policies. Sensitive information is information derived from sources that are confidential or private. Any IFC system has two components—the *policy component* and the *enforcement component*. The policy component allows labeling of private information sources. The label on a source specifies how private information from that source can be used and where it can flow. The collection of rules for labeling is called the policy. The enforcement component enforces policies. WebPol contributes a policy component to complement existing work on enforcement components in web browsers. Many existing enforcement components can be used with WebPol. For completeness, we describe both policy and enforcement components here.

Policy component. The policy component provides a way to label or mark sensitive data sources with labels that represent confidentiality and where data can flow. In the context of webpages, data sources are objects generated in response to user events like the content of a password box generated due to key presses or a mouse click on a sensitive button, and data obtained in a network receive event. In WebPol, data sources can be labeled with three kinds of labels, in increasing order of confidentiality: 1) the label `public` represents non-sensitive data, 2) for each domain `domain`, the label `domain` represents data private to the domain; such data’s flow should be limited only to the browser and servers belonging to `domain` and its subdomains, and 3) the label `local` represents very confidential data that must never leave the browser. Technically, labels are elements of the partial order $\text{public} < \text{domain}_i < \text{local}$. Labels higher in the order represent more confidentiality than labels lower in the order. These labels are fairly expressive.¹ For example, labeling a data source with the domain of the hosting page prevents exfiltration to third-parties. Labeling a data source with the domain of a third-party provider such as an page analytics provider allows transfer to only that service.

Since most data on a webpage is not sensitive, it is reasonable to label data sources `public` by default and only selectively assign a different label. WebPol

¹ Richer label models that support, for instance, conjunctions and disjunctions of labels [27] are compatible with WebPol. However, we have not found the need for such models so far.

uses this blacklisting approach. Two nuances of source labeling are noteworthy. The first is its fine granularity. Not all objects generated by the same class of events have the same label. For instance, characters entered in a password field may have the domain label of the hosting page, limiting their flow only to the host, but characters entered in other fields may be accessible to third-party advertising or analytics scripts without restrictions. This leads to the following requirement on the policy component.

Requirement 1: The policy component must allow associating different policies with different elements of the page.

The second nuance is that the label of an object can be dynamic, i.e., history-dependent. Consider a policy that hides from an analytics script how many times a user clicked within an interactive panel, but wants to share whether or not the user clicked at least once. The label of a click event on the panel is `public` the first time the user clicks on it and `private` afterwards and, hence, it depends on the history of user interaction. This yields the following requirement on the policy component.

Requirement 2: Labels may be determined dynamically. This requirement means that labels must be set by *trusted policy code* that is executed on-the-fly and that has local state.

Enforcement component. Source data labels must be enforced even as scripts transform and transmit the data. Existing literature is rife with techniques for doing this, even in the context of web-browsers. Fine-grained taint tracking [18,16,7,26,9], coarse-grained taint tracking [28,6], secure multi-execution [13,30], faceted execution [4,32] and static analysis [15,10,21] are some enforcement techniques that have been considered in the context of JavaScript. They differ considerably in their mechanics, their expressiveness and ease of fit with the browser programming model. WebPol has been designed keeping fine-grained taint tracking (FGTT) in mind, so we explain that technique in some detail below.

In FGTT, the language runtime is modified to track information flows and to attach a label (often called a taint) with each runtime object, including objects on the stack, the heap and, in the context of web browsers, the DOM. Two kinds of flows are typically considered. *Explicit* flows arise as a result of direct assignment. In these cases, the label of the destination object is overwritten with the label of the source object. *Implicit* flows arise due to control dependencies. For instance, in `pub = false; if (sec) pub = true`, the final value of `pub` depends on `sec` although there is no direct assignment from `sec` to `pub`. Implicit flows are tracked by keeping a *context* label on the instruction pointer. Once both explicit and implicit flows are tracked, enforcing policies is straightforward: An outgoing communication with domain `d`'s servers is allowed only if the labels on the payload of the communication and the instruction pointer at the point of the communication are either `public` or `d`. This ensures that all labels attached to source data are respected.

FGTT can be implemented either by modifying the browser's JavaScript engine to track flows and labels [16,7], or by a source-to-source transform of

JavaScript code prior to execution [9]. There is a space and time overhead associated with storing labels and tracking them. However, with careful engineering, this overhead can be reduced enough to not be noticeable to end-users.

3 WebPol policy model

WebPol works on a browser that has already been augmented with IFC enforcement. It provides a framework that allows setting labels at fine-granularity, thus expressing and enforcing rich policies. This section describes the threat model for WebPol and explains the WebPol design.

Threat model. WebPol prevents under-the-hood exfiltration of sensitive data that has been provided to third-party scripts for legitimate reasons. So, third-party scripts are not trusted but code from the host domain is trusted.

We are interested only in JavaScript-level bugs or exfiltration attempts. We trust the browser infrastructure to execute all JavaScript code following the language’s semantics and to dispatch events correctly. Low-level attacks that target vulnerabilities in the browser engine are out of scope. Similarly, defending against network attacks (like man-in-the-middle attacks) is not our goal. Orthogonal techniques like end-to-end encryption or HTTPS can be used to defend against those attacks. Integrity attacks are also out of scope. For instance, attacks based on sending requests containing no sensitive data to websites, where the user might already be logged in, cannot be prevented using this model.

WebPol executes on top of an IFC enforcement in the browser. That enforcement is assumed to be correct and to track all flows. Prior work on such enforcement has often been supplemented with formal proofs to show that the enforcement is correct, at least abstractly [16,7,26].

WebPol’s policies are agnostic to specific channels of information leak. However, current IFC enforcements in browsers track only explicit and implicit flows. Consequently, leaks over other channels such as timing and memory-usage are currently out of scope. As IFC enforcements improve to cover more channels, WebPol’s policies will extend to them as well.

3.1 Policies as event handlers

The first question in the design of WebPol is who should specify policies. Since our goal is to prevent exfiltration of data by third-party scripts and it is the developer of the host page who bootstraps the inclusion of scripts and best understands how data on the page should be used, it is natural and pragmatic to have the developer specify policies, possibly as part of the page itself.

The next question is how the developer specifies policies. To answer this, we recall the two requirements we identified in Section 2—it should be possible to specify different policies on different page elements and policies should be allowed to include code that is executed on-the-fly to generate labels. When we also consider the fact that sensitive data is usually generated by input events, it is clear that policies should *be* page element-specific, (trusted) code that is executed

after events have occurred (this code labels event-generated data). Fortunately, web browsers provide exactly this abstraction in the form of event handlers! So, we simply extend the event-handling logic in web browsers to express WebPol policies. This allows us to leverage a lot of the existing browser logic for event handler installation, parsing and event dispatch. Before explaining how we do this, we provide a brief overview of event handling in web browsers.

Event handlers and event dispatch. Browsers execute JavaScript functions, called event handlers, in response to input events like mouse clicks, key presses, and asynchronous network receives. Save for network receive events, every event has a *target*, which is an element in the page’s DOM where the event originated. For instance, if a button is clicked, the target of the ensuing event is the button. Code running on a page can add an event handler on any element on the page, listening for a specific event. When an event occurs, all handlers associated for that event with the event’s target and the target’s ancestors are triggered sequentially. This is called *event dispatch*. The specific order in which handlers are triggered is not relevant for our purposes (although it is fairly interesting for IFC enforcement [26]). The whole process is bootstrapped by the static HTML of the page, which may contain JavaScript that is executed when the page loads initially, and this JavaScript installs the first set of event handlers.

Policy handlers. In WebPol, policies are special event handlers, specified using a special marker in the HTML source of the hosting page. These special handlers, called *policy handlers*, follow standard JavaScript syntax, can be attached to any page element, listening for any event and, like other handlers, are triggered every time the event is dispatched on the element or any of its descendants in the DOM. However, unlike other handlers, the sole goal of policy handlers is to assign labels to other sensitive objects, including the event being dispatched. To allow the policy handlers to do this, we modify the browser slightly to afford these handlers two special privileges:

- Policy handlers can execute two new JavaScript API functions that set labels on other objects. No other JavaScript code can execute these two functions. These functions are described later.
- During event dispatch all applicable policy handlers are executed before ordinary handlers. This ensures that labels are set before ordinary handlers (including those of third-party scripts) execute.

To maintain the integrity of the policies, policy handlers must be included in the HTML source of the page directly. They *cannot* be installed dynamically by JavaScript code. Otherwise, third-party scripts could install policy handlers that set very permissive labels. Also, if a DOM element has a policy handler, we disallow third-party scripts from detaching that element or moving it elsewhere, as that can change the interpretation of the policy. Similarly, changing the attributes of such an element is restricted.

Since different policy handlers can be associated with different elements, Requirement 1 is satisfied. Moreover, policy handlers are ordinary JavaScript

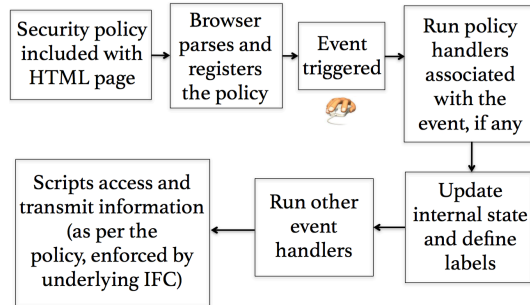


Fig. 1: Workflow of the WebPol policy model

code, so they can also maintain local state in private variables, thus satisfying Requirement 2.

The workflow of policy interpretation in WebPol is shown in Figure 1. We briefly summarize the steps:

1. The web page developer specifies the policy in the host HTML page in the form of special event handlers.
2. The browser parses the policy and registers its handlers (mostly like usual handlers, but with the two special privileges mentioned above).
3. When an event dispatches, listening policy handlers are executed first.
4. These policy handlers set labels on objects affected by the event, including the event object itself. They may also update any local state they maintain.
5. The remaining event handlers are dispatched as usual. The IFC enforcement in the browser enforces all labels that have been set by the policy handlers (during any prior event’s dispatch), thus preventing any data leak in contravention of the labels.

3.2 Integration with the web browser

WebPol needs minor modifications to the browser to parse and interpret policies and to expose additional JavaScript API functions to set labels.

HTML and event dispatch changes. WebPol adds an HTML extension to differentiate policy code from other JavaScript code. Concretely, we change the browser’s parser to interpret any script file with the extension `.policy` included directly in the host page as a policy. If such a *policy script* installs a handler, it is treated as a policy handler. Additionally, a policy script can set labels on the page’s global variables and DOM elements (like password fields). If a script does this, it should be included in the host page before third-party scripts that use those variables. WebPol also requires a small change to the browser’s event dispatch mechanism to execute policy handlers before other handlers.

```

1  var p = document.getElementById("pwd");
2  p.addEventListener("keypress", function (e){
3    var score = checkPwdStrength(p.value);
4    document.getElementById("pwdStrength").innerText = score;
5    new Image().src = "http://stealer.com/pwd.jsp?pwd="+p +score;
6  });

```

Listing 1: Password strength checking script that leaks the password

Label-setting APIs. WebPol exposes two new JavaScript API functions to set labels. These functions can be called only by the policy code in `.policy` files and handlers installed by such files (we modify the browser to enforce this).

The function `setLabel(label)` sets the label of the object on which it is called to `label`. As explained earlier, `label` can be `public`, a domain name, or `local` (the default is `public`). Once an object’s label is set, it is enforced by the underlying IFC enforcement. The special label `HOST` is a proxy for the domain of the host page.

The function `setContext(label)` can be called only on an event object. It restricts the visibility of the event to label `label` and higher. In simple terms, if `label` is a domain, then only that domain can ever learn that this event occurred, whereas if `label` is `local`, then no domain can ever learn that this event occurred. Technically, this is accomplished by setting the so-called *pc* or program counter label of event handlers running during the dispatch to `label`, which ensures that their side-effects (writes to DOM and network communication) are labeled `label` or higher.

As opposed to `setLabel`, which makes individual data objects (like password fields) private, `setContext` makes the *existence* of an event private. This is useful. For instance, clicking on the “politics” section of a news feed might indicate that the user is interested in politics, which may be private information, so the page may want to hide even the existence of click events from third-party scripts. (The distinction between the privacy of event content and event occurrence has been previously described by Rafnsson and Sabelfeld [25].)

4 Examples

We illustrate the expressiveness of WebPol policies through a few examples.

Example 1: Password strength checker. Many websites deploy *password strength checkers* on pages where users set new passwords. A password strength checker is an event handler from a third-party library that is triggered each time the user enters a character in the new password field. The handler provides visual feedback to the user about the strength of the password entered so far. Strength checkers usually check the length of the password and the diversity of characters used. Consequently, they do not require any network communication. However, standard browser policies cannot enforce this and the password strength checker


```

1 function currencyConverter() {
2   var toCur = document.getElementById("to").value;
3   var xh = new XMLHttpRequest();
4   xh.onreadystatechange = function() {
5     if (xh.readyState == 4) {
6       currencyRate = eval(xhttp.responseText);
7       var aAmt = document.getElementById("amt").value;
8       var convAmt = aAmt * currencyRate;
9       document.getElementById("camt").innerHTML = convAmt;
10      xh.open("GET","http://currConv.com/amount.jsp?atc=" + aAmt);
11      xh.send(); }
12  xh.open("GET","http://currConv.com/conv.jsp?toCur=" + toCur, true);
13  xh.send(); }

```

Listing 2: Currency converter script that leaks a private amount

can easily leak the password if it wants to. Listing 1 shows such a “leaky” password checker. The checker installs a listener for keypresses in the password field (line 2). In response to every keypress, the listener delivers its expected functionality by checking the strength of the password and indicating this to the user (lines 3, 4), but then it leaks out the password to `stealer.com` by requesting an image at a url that includes the password (line 5). The browser’s standard SOP allows this.

With WebPol, the developer of the host webpage can prevent any exfiltration of the password by including the policy script:

```
document.getElementById("pwd").setLabel("HOST");
```

This policy sets the label of the password field to the host’s own domain using the function `setLabel()`. Subsequently, the IFC enforcement restricts all outgoing communication that depends on the password field to the host.

Conceptually, this example is simple because it does not really leverage the fine-granularity of WebPol policies and FGTT. Here, the third-party script does not need any network communication for its intended functionality and, hence, simpler confinement mechanisms that prohibit a third-party script from communicating with remote servers would also suffice. Our next example is a scenario where the third-party script legitimately needs remote communication. It leverages the fine-granularity of WebPol policies and FGTT.

Example 2: Currency conversion. Consider a webpage from an e-commerce website which displays the cost of an item that the user intends to buy. The amount is listed in the site’s native currency, say US dollars (USD), but for the user’s convenience, the site also allows the user to see the amount converted to a currency of his/her choice. For this, the user selects a currency from a dropdown list. A third-party JavaScript library reads both the USD amount and the second currency, converts the amount to the second currency and inserts it into the webpage, next to the USD amount. The third-party script fetches the current conversion rate from its backend service at `currConv.com`. Consequently,

```

1 var p = document.getElementById("sect_name");
2 p.addEventListener("click",function(event){
3   event.setLabel("HOST"); });

```

Listing 3: Policy that allows counting clicks but hides details of the clicks

```

1 clickCount = 0;
2 var p = document.getElementById("sect_name");
3 p.addEventListener("click",function clkHdlr(e){ clickCount += 1; });

```

Listing 4: Analytics script that counts clicks

it must send the *name* of the second currency to its backend service, but must not send the amount being converted (the amount is private information). The web browser’s same-origin policy has been relaxed (using, say, CORS [20]) to allow the script to talk to its backend service at `currConv.com`. The risk is that the script can now exfiltrate the private amount. Listing 2 shows a leaky script that does this. On line 13, the script makes a request to its backend service passing to the second currency. The callback handler (lines 4–11) reads the amount from the page element `amt`, converts it and inserts the result into the page (lines 6–9). Later, it leaks out the amount to the backend service on line 10, in contravention of the intended policy.

With WebPol, this leak can be prevented with the following policy that sets the label of the amount to the host only:

```
document.getElementById("amt").setLabel("HOST")
```

This policy prevents exfiltration of the amount but does not interfere with the requirement to exfiltrate the second currency. Importantly, no modifications are required to a script that does not try to leak data (e.g., the script obtained by dropping the leaky line 10 of Listing 2).

Example 3: Web analytics. To better understand how users interact with their websites, web developers often include third-party analytics scripts that track user clicks and keypresses to generate useful artifacts like page heat-maps (which part of the page did the user interact with the most?). Although a web developer might be interested in tracking only certain aspects of their users’ interaction, the inclusion of the third-party scripts comes with the risk that the scripts will also record and exfiltrate other private user behavior (possibly for monetizing it later). Using WebPol, the web developer can write precise policies about which user events an analytics script can access and when. We show several examples of this.

To allow a script to only count the number of occurrences of a class of events (e.g., mouse clicks) on a section of the page, but to hide the details of the individual events (e.g., the coordinates of every individual click), the web developer can add a policy handler on the top-most element of the section to set the label of the individual event objects to `HOST`. This prevents the analytics script’s listening handler from examining the details of individual events, but since the

```

1  var alreadyClicked = false;
2  var p = document.getElementById("sect_name");
3  p.addEventListener("click",function(event){
4    if (alreadyClicked = true) event.setContext("HOST");
5    else {alreadyClicked = true; event.setLabel("HOST");} });

```

Listing 5: Policy that tracks whether a click happened or not only

handler is still invoked at each event, it can count their total number. Listings 3 and 4 show the policy handler and the corresponding analytics script that counts clicks in a page section named `sect_name`.

Next, consider a restriction of this policy, which allows the analytics script to learn only whether or not *at least one* click happened in the page section, completely hiding clicks beyond the first. This policy can be represented in WebPol using a local state variable in the policy to track whether or not a click has happened. Listing 5 shows the policy. The policy uses a variable `alreadyClicked` to track whether or not the user has clicked in the section. Upon the user's first click, the policy handler sets the event's label to the host's domain (line 5). This makes the event object private but allows the analytics handler to trigger and record the occurrence of the event. On every subsequent click, the policy handler sets the event's *context* to the host domain using `setContext()` (line 4). This prevents the analytics script from exfiltrating any information about the event, including the fact that it occurred.

Finally, note that a developer can subject different page sections to different policies by attaching different policy handlers to them. The most sensitive sections may have a policy that unconditionally sets the event context to the host's, effectively hiding all user events in those sections. Less sensitive sections may have policies like those of Listings 5 and 3. Non-sensitive sections may have no policies at all, allowing analytics scripts to see all events in them.

Example 4: Defending against overlay-based attacks. The full version of the paper [8] describes a simple WebPol policy that defends against an attack where a malicious script creates a transparent overlay over a sensitive element (like a password field) to record user events like keypresses without policy protection.

5 Implementation

We have prototyped WebPol in WebKit, a popular open source browser engine that powers many browsers including Apple's Safari. Our implementation runs on top of our prior IFC enforcement in WebKit that uses FGTT and a bit of on-the-fly static analysis [7]. The IFC enforcement is highly optimized, and covers most JavaScript native functions (the DOM API) [26]. It targets WebKit nightly build `#r122160` and works with the Safari web browser, version 6.0. Since it is difficult to port our earlier implementation (not WebPol) to a newer version of WebKit, we choose to evaluate WebPol on this slightly outdated setup. This

suffices since WebPol’s design is not affected by recent browser updates. The source code is available online at: <https://github.com/bichhawat/ifc4bc>.

Our earlier IFC implementation modified approximately 6,800 lines in the JavaScript engine, the DOM APIs and the event handling logic for FGTT. To implement WebPol, we additionally modified the HTML parser to distinguish policy files (extension `.policy`) from other JavaScript files and to give policy code extra privileges. We also added the two new JavaScript API functions `setLabel()` and `setContext()`. Finally, we modified the event dispatch logic to trigger policy handlers before other handlers. In all, we changed 25 lines in the code of the parser, added 60 lines for the two new API functions and changed 110 lines in the event dispatch logic. Overall, implementing WebPol has low overhead, and we expect that it can also be ported to other browsers or later versions of WebKit easily.

6 Evaluation

The goal of our evaluation is two-fold. First, we want to measure WebPol overhead, both on parsing and installing policies during page load and on executing policy handlers later. We do this for four examples presented in Section 4 and for two real-world websites. Second, we wish to understand whether WebPol can be used easily. Accordingly, we apply WebPol policies to two real-world websites and report on our experience. All our experiments were performed on a 3.2GHz Quad-core Intel Xeon processor with 8GB RAM, running Mac OS X version 10.7.4 using the browser configuration described in Section 5.

Performance overheads on synthetic examples. To measure WebPol’s runtime overhead, we tested four examples from Section 4 (Examples 1, 2 and the two sub-examples of Example 3) in three different configurations: **Base**—uninstrumented browser, no enforcement; **IFC**—taint tracking from prior work, but no policy handlers (everything is labeled public); **WebPol**—our system running policy handlers and taint tracking.

JavaScript execution time: To measure the overheads of executing policy handler code, we interacted with all four programs manually by entering relevant data and performing clicks a fixed number of times. For each of these configurations, we measured the total time spent *only in executing JavaScript*, including scripts and policies loaded initially with the page and the scripts and policies executed in response to events. The difference between **IFC** and **Base** is the overhead of taint tracking, while the difference between the **WebPol** and **IFC** is the overhead of evaluating policy handlers. Since we are only measuring JavaScript execution time and there are no time-triggered handlers in these examples, variability in the inter-event gap introduced by the human actor does not affect the measurements.

The left half of Table 1 shows our observations. All numbers are averages of 5 runs and the standard deviations are all below 7%. Taint-tracking (**IFC**) adds overheads ranging from 15.5% to 26.7% over **Base**. To this, policy handlers (**WebPol**) adds overheads ranging from 1.9% to 29.2%. **WebPol** overheads are

Example #	JavaScript Execution Time			Page Load Time		
	Base	IFC	WebPol	Base	IFC	WebPol
Example 1	2430	2918 (+20.1%)	2989 (+1.9%)	16	17 (+6.3%)	19 (+12.5%)
Example 2	3443	4361 (+26.7%)	5368 (+29.2%)	41	43 (+4.9%)	46 (+7.2%)
Example 3 (count)	1504	1737 (+15.5%)	1911 (+11.6%)	24	25 (+4.2%)	31 (+25.0%)
Example 3 (presence)	1780	2095 (+17.7%)	2414 (+18.9%)	26	28 (+7.7%)	30 (+7.7%)

Table 1: Performance of examples from Section 4. All time in ms. The percentages in parentheses in the column **IFC** are overheads relative to **Base**. Similar numbers in the column **WebPol** are *additional* overheads, still relative to **Base**.

already modest, but we also note that this is also a very challenging (conservative) experiment for WebPol. The scripts in both sub-examples of Example 3 do almost nothing. The scripts in Examples 1 and Example 2 are slightly longer, but are still much simpler than real scripts. On real and longer scripts, the relative overheads of evaluating the policy handlers is significantly lower as shown later. Moreover, our baseline in this experiment does not include other browser costs, such as the cost of page parsing and rendering, and network delays. Compared to those, both **IFC** and **WebPol** overheads are negligible.

Page load time: We separately measured the time taken for loading the initial page (up to the DOMContentLoaded event). The difference between **WebPol** and **IFC** is the overhead for parsing and loading policies. The right half of Table 1 shows our observations. All numbers are the average of 20 runs and all standard deviations are below 8%. WebPol overheads due to policy parsing and loading range from 7.2% to 25% (last column). When we add overheads due to taint tracking (column **IFC**), the numbers increase to 12.1% to 29.2%. Note that page-load overheads are incurred only once on every page (re-)load.

Real-world websites. To understand whether WebPol scales to real-world websites, we evaluated WebPol on policies for two real-world applications—the website <http://www.passwordmeter.com> that deploys a password-strength checker (similar to Example 1) and a bank login page that includes third-party analytics scripts (similar to Example 3). Both policies were written by hand and are shown in the full version of the paper [8].

Experience writing policies: In both cases, we were able to come up with meaningful policies easily after we understood the code, suggesting that WebPol policies can be (and should be) written by website developers. The policy for the password-strength checker is similar to Listing 1 and prevents the password from being leaked to third-parties. We had to write four lines of additional policy code to allow the script to write the results of the password strength check (which depends on the password) into the host page. The analytics script on the bank

Website	JavaScript Execution Time			Page Load Time		
	Base	IFC	WebPol	Base	IFC	WebPol
Password	79.5	115.5 (+45.3%)	126 (+13.2%)	303	429 (+41.6%)	441 (+4.0%)
Analytics	273.4	375.1 (+37.2%)	386.1 (+4.0%)	2151	2422 (+12.6%)	2499 (+3.6%)

Table 2: Performance on two real-world websites. All time in ms. The percentages in parentheses in the column **IFC** are overheads relative to **Base**. Similar numbers in the column **WebPol** are *additional* overheads, still relative to **Base**.

website communicates all user-behavior to its server. We specified a policy that disallows exfiltration of keypresses on the username and the password text-boxes to third-parties.

Performance overheads: We also measured performance overheads on the two websites, in the same configurations as for the synthetic examples. Table 2 shows the results. On real-world websites, where actual computation is long, the overheads of WebPol are rather small. The overheads of executing policy handlers, relative to **Base**’s JavaScript execution time, are 4.0% and 13.2%, while the overheads of parsing and loading policies are no more than 4.0%. Even the total overhead of **IFC** and **WebPol** does not adversely affect the user experience in any significant way.

This experiment indicates that WebPol is suitable for real-world websites.

7 Related Work

Browser security is a very widely-studied topic. Here, we describe only closely related work on browser security policies and policy enforcement techniques.

Information flow control and script isolation. The work most closely related to our is that of Vanhoef *et al.* [30] on stateful declassification policies in reactive systems, including web browsers. Their policies are similar to ours, but there are significant differences. First, their policies are attached to the browser and they are managed by the browser user rather than website developers. Second, the policies have coarse-granularity: They apply uniformly to all events of a certain type. Hence, it is impossible to specify a policy that makes keypresses in a password field secret, but makes other keypresses public. Third, the enforcement is based on secure multi-execution [13], which is, so far, not compatible with shared state like the DOM.

COWL [28] enforces mandatory access control at coarse-granularity. In COWL, third-party scripts are sandboxed. Each script gets access to either remote servers or the host’s DOM, but not both. Scripts that need both must be re-factored to pass DOM elements over a message-passing API (`postMessage`). This can be both difficult and have high overhead. For scripts that do not need this factorization, COWL is more efficient than solutions based on FGTT.

Mash-IF [21] uses static analysis to enforce IFC policies. Mash-IF’s model is different from WebPol’s model. Mash-IF policies are attached only to DOM nodes and there is no support for adding policies to new objects or events. Also, in Mash-IF, the browser user (not the website developer) decides what declassifications are allowed. Mash-IF is limited to a JavaScript subset that excludes commonly used features such as `eval` and dynamic property access.

JSand [3] uses server-side changes to the host page to introduce wrappers around sensitive objects, in the style of object capabilities [24]. These wrappers mediate every access by third-party scripts and can enforce rich access policies. Through secure multi-execution, coarse-grained information flow policies are also supported. However, as mentioned earlier, it is unclear how secure multi-execution can be used with scripts that share state with the host page.

WebPol policies are enforced using an underlying IFC component. Although, in principle, any IFC technique such as fine-grained taint tracking [18,16,7,9], coarse-grained taint tracking [28] or secure multi-execution [13] can be used with WebPol, to leverage the full expressiveness of WebPol’s finely-granular policies, a fine-grained IFC technique is needed. JSFlow [17,16] is a stand-alone implementation of a JavaScript interpreter with fine-grained taint tracking. Many seminal ideas for labeling and tracking flows in JavaScript owe their lineage to JSFlow, but since JSFlow is written from scratch it has very high overheads. Building on ideas introduced by Just *et al.* [19], our own prior work [26,7] implements fine-grained IFC in an existing browser engine, WebKit, by modifying the JavaScript interpreter. The overheads are significantly lower than JSFlow, which is why chose to integrate WebPol with our own work. Both JSFlow and our work include formal proofs that the taint tracking is complete, relative to the abstractions of a formal model. Chudnov and Naumann [9] present another approach to fine-grained IFC for JavaScript. They rewrite source programs to add shadow variables that hold labels and additional code that tracks taints. This approach is inherently more portable than that of JSFlow or our work, both of which are tied to specific, instrumented browsers. However, it is unclear to us how this approach could be extended with a policy framework like WebPol that assigns state-dependent labels at runtime.

Access control. The traditional browser security model is based on restricting scripts’ access to data, not on tracking how scripts use data. However, no model based on access control alone can simultaneously allow scripts access to data they need for legitimate purposes and prevent them from leaking the data on the side. Doing so is the goal of IFC and WebPol. Nonetheless, we discuss some related work on access control in web browsers.

The standard same-origin policy (SOP) and content-security policy (CSP) were described in Section 1. An additional, common access policy—cross-origin resource sharing (CORS) [20]—relaxes SOP to allow some cross-origin requests.

Conscript [23] allows the specification of fine-grained access policies on individual scripts, limiting what actions every script can perform. Similarly, Ad-Jail [22] limits the execution of third-party scripts to a shadow page and restricts communication between the script and the host page. Zhou and Evans [33] take a

dual approach, where fine-grained access control rules are attached to DOM elements. The rules specify which scripts can and cannot access individual elements. Along similar lines, Dong *et al.* [14] present a technique to isolate sensitive data using authenticated encryption. Their goal is to reduce the size of the trusted computing base. ADsafe [11] and FBJS [1] restrict third-party code to subsets of JavaScript, and use static analysis to check for illegitimate access. Caja [2] uses object capabilities to mediate all access by third-party scripts. WebJail [29] supports least privilege integration of third-party scripts by restricting script access based on high-level policies specified by the developer. All these techniques enforce only access policies and cannot control what a script does with data it has been provided in good faith.

8 Conclusion

Third-party JavaScript often requires access to sensitive data to provide meaningful functionality, but comes with the risk that the data may be leaked on the side. Information flow control in web browsers can solve this problem. Within this context, this paper proposed WebPol, a mechanism for labeling sensitive data, dynamically and at fine-granularity. WebPol uses JavaScript for policy specification, which makes it developer-friendly, and re-uses the browser's event handling logic for policy interpretation, which makes it easy to implement and improves the likelihood of easy portability across browsers and versions. Our evaluation indicates that WebPol has low-to-moderate overhead, even including the cost of information flow control and, hence, it can be used on websites today.

Acknowledgments. We thank several anonymous reviewers for their excellent feedback. This work was funded in part by the Deutsche Forschungsgemeinschaft (DFG) grant “Information Flow Control for Browser Clients” under the priority program “Reliably Secure Software Systems” (RS³).

References

1. Facebook. FBJS. <https://developers.facebook.com/docs/javascript>, [Online; accessed 19-Jun-2017]
2. Google Caja - A source-to-source translator for securing JavaScript-based web content. <https://developers.google.com/caja/>, [Online; accessed 19-Jun-2017]
3. Agten, P., Van Acker, S., Brondsema, Y., Phung, P.H., Desmet, L., Piessens, F.: JSand: Complete Client-side Sandboxing of Third-party JavaScript Without Browser Modifications. In: Proc. 28th Annual Computer Security Applications Conference (ACSAC). pp. 1–10. ACM, New York, NY, USA (2012)
4. Austin, T.H., Flanagan, C.: Multiple Facets for Dynamic Information Flow. In: Proc. 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL). pp. 165–178. ACM, New York, NY, USA (2012)
5. Barth, A.: The Web Origin Concept. <http://tools.ietf.org/html/rfc6454>, [Online; accessed 19-Jun-2017]

6. Bauer, L., Cai, S., Jia, L., Passaro, T., Stroucken, M., Tian, Y.: Run-time Monitoring and Formal Analysis of Information Flows in Chromium. In: Proc. ISOC Network and Distributed System Security Symposium (NDSS) (2015)
7. Bichhawat, A., Rajani, V., Garg, D., Hammer, C.: Information Flow Control in WebKit's JavaScript Bytecode. In: Proc. 3rd International Conference on Principles of Security and Trust (POST), LNCS, Vol. 8414, pp. 159–178. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
8. Bichhawat, A., Rajani, V., Jain, J., Garg, D., Hammer, C.: WebPol: Fine-grained Information Flow Policies for Web Browsers (Full version). <http://arxiv.org/abs/1706.06932> (2017)
9. Chudnov, A., Naumann, D.A.: Inlined Information Flow Monitoring for JavaScript. In: Proc. 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS). pp. 629–643. ACM, New York, NY, USA (2015)
10. Chugh, R., Meister, J.A., Jhala, R., Lerner, S.: Staged Information Flow for JavaScript. In: Proc. 30th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI). pp. 50–62. ACM, New York, NY, USA (2009)
11. Crockford, D.: ADsafe. <http://adsafe.org/>, [Online; accessed 19-Jun-2017]
12. De Groef, W., Devriese, D., Nikiforakis, N., Piessens, F.: FlowFox: A Web Browser with Flexible and Precise Information Flow Control. In: Proc. 19th ACM Conference on Computer and Communications Security (CCS). pp. 748–759. ACM, New York, NY, USA (2012)
13. Devriese, D., Piessens, F.: Noninterference through Secure Multi-execution. In: Proc. 31st IEEE Symposium on Security and Privacy (SP). pp. 109–124. IEEE Computer Society, Washington, DC, USA (2010)
14. Dong, X., Chen, Z., Siadati, H., Tople, S., Saxena, P., Liang, Z.: Protecting Sensitive Web Content from Client-side Vulnerabilities with CRYPTONS. In: Proc. 20th ACM SIGSAC Conference on Computer and Communications Security (CCS). pp. 1311–1324. ACM, New York, NY, USA (2013)
15. Guarneri, S., Pistoia, M., Tripp, O., Dolby, J., Teilhet, S., Berg, R.: Saving the World Wide Web from Vulnerable JavaScript. In: Proc. 2011 International Symposium on Software Testing and Analysis (ISSTA). pp. 177–187. ACM, New York, NY, USA (2011)
16. Hedin, D., Birgisson, A., Bello, L., Sabelfeld, A.: JSFlow: Tracking Information Flow in JavaScript and Its APIs. In: Proc. 29th Annual ACM Symposium on Applied Computing (SAC). pp. 1663–1671. ACM, New York, NY, USA (2014)
17. Hedin, D., Sabelfeld, A.: Information-Flow Security for a Core of JavaScript. In: Proc. IEEE 25th Computer Security Foundations Symposium (CSF). pp. 3–18. IEEE Computer Society, Washington, DC, USA (2012)
18. Jang, D., Jhala, R., Lerner, S., Shacham, H.: An Empirical Study of Privacy-violating Information Flows in JavaScript Web Applications. In: Proc. 17th ACM Conference on Computer and Communications Security (CCS). pp. 270–283. ACM, New York, NY, USA (2010)
19. Just, S., Cleary, A., Shirley, B., Hammer, C.: Information Flow Analysis for JavaScript. In: Proc. 1st ACM SIGPLAN International Workshop on Programming Language and Systems Technologies for Internet Clients (PLASTIC). pp. 9–18. ACM, New York, NY, USA (2011)
20. van Kesteren, A.: Cross-Origin Resource Sharing. <http://www.w3.org/TR/cors/>, [Online; accessed 19-Jun-2017]
21. Li, Z., Zhang, K., Wang, X.: Mash-IF: Practical information-flow control within client-side mashups. In: Proc. 40th Annual IEEE/IFIP International Conference on Dependable Systems Networks (DSN). pp. 251–260 (2010)

22. Louw, M.T., Ganesh, K.T., Venkatakrisnan, V.N.: AdJail: Practical Enforcement of Confidentiality and Integrity Policies on Web Advertisements. In: Proc. 19th USENIX Conference on Security (USENIX Security). pp. 24–40. USENIX Association, Berkeley, CA, USA (2010)
23. Meyerovich, L.A., Livshits, B.: ConScript: Specifying and Enforcing Fine-Grained Security Policies for JavaScript in the Browser. In: Proc. 31st IEEE Symposium on Security and Privacy (SP). pp. 481–496. IEEE Computer Society, Washington, DC, USA (2010)
24. Miller, M.: Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control. Ph.D. thesis, Johns Hopkins University (2006)
25. Rafnsson, W., Sabelfeld, A.: Secure Multi-execution: Fine-Grained, Declassification-Aware, and Transparent. In: Proc. IEEE 26th Computer Security Foundations Symposium (CSF). pp. 33–48. IEEE Computer Society, Washington, DC, USA (2013)
26. Rajani, V., Bichhawat, A., Garg, D., Hammer, C.: Information Flow Control for Event Handling and the DOM in Web Browsers. In: Proc. IEEE 28th Computer Security Foundations Symposium (CSF). pp. 366–379. IEEE Computer Society, Washington, DC, USA (2015)
27. Stefan, D., Russo, A., Mazières, D., Mitchell, J.C.: Disjunction Category Labels. In: Proc. 16th Nordic Conference on Information Security Technology for Applications (NordSec). pp. 223–239. Springer-Verlag, Berlin, Heidelberg (2011)
28. Stefan, D., Yang, E.Z., Marchenko, P., Russo, A., Herman, D., Karp, B., Mazières, D.: Protecting Users by Confining JavaScript with COWL. In: Proc. 11th USENIX Conference on Operating Systems Design and Implementation (OSDI). pp. 131–146. USENIX Association, Berkeley, CA, USA (2014)
29. Van Acker, S., De Ryck, P., Desmet, L., Piessens, F., Joosen, W.: WebJail: Least-privilege Integration of Third-party Components in Web Mashups. In: Proc. 27th Annual Computer Security Applications Conference (ACSAC). pp. 307–316. ACM, New York, NY, USA (2011)
30. Vanhoef, M., De Groef, W., Devriese, D., Piessens, F., Rezk, T.: Stateful Declassification Policies for Event-Driven Programs. In: Proc. IEEE 27th Computer Security Foundations Symposium (CSF). pp. 293–307. IEEE Computer Society, Washington, DC, USA (2014)
31. West, M.: Content Security Policy Level 3. <https://www.w3.org/TR/CSP3/>, [Online; accessed 19-Jun-2017]
32. Yang, J., Yessenov, K., Solar-Lezama, A.: A Language for Automatically Enforcing Privacy Policies. In: Proc. 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL). pp. 85–96. ACM, New York, NY, USA (2012)
33. Zhou, Y., Evans, D.: Protecting Private Web Content from Embedded Scripts. In: Atluri, V., Diaz, C. (eds.) Proc. 16th European Symposium on Research in Computer Security (ESORICS), LNCS, Vol. 6879, pp. 60–79. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)