

# HTTPi for Practical End-to-End Web Content Integrity

Kapil Singh  
Georgia Institute of Technology

Helen J. Wang  
Microsoft Research

Alexander Moshchuk  
Microsoft Research

Collin Jackson  
Carnegie Mellon University

Wenke Lee  
Georgia Institute of Technology

## ABSTRACT

The widespread growth of open wireless hot spots has made it very easy for network attackers to carry out man-in-the-middle attacks and impersonate web sites. End-to-end security between a user’s web browser and web sites is ever more needed to allow meaningful enforcement of the same-origin policy on the web browser platform. Although HTTPS can be used to prevent such attacks, its universal adoption by web sites is hindered by its performance cost and its inability to be cached at intermediate servers (such as CDN servers and cache proxies) while maintaining end-to-end security. With significant and increasing amount of web content being cacheable, HTTPS is not the complete answer to an end-to-end secure web.

In this paper, we observe that only end-to-end authentication and integrity are required for the browser platform to meaningfully enforce the same-origin policy. Without end-to-end confidentiality, content can be cached. In light of this observation, we propose a new protocol, HTTPi, which offers only end-to-end authentication and integrity. HTTPi works seamlessly with and benefits from the existing web caching infrastructure. It performs content signing while preserving progressive content loading supported by browsers. Because content signing can be done offline, HTTPi incurs negligible overhead over HTTP. We advocate that sites use HTTPS for requests that require end-to-end confidentiality, and HTTPi for all other requests. Our prototype and evaluation experience show that HTTPi is practical for adoption.

## 1. INTRODUCTION

The same-origin policy [33] (SOP) is the key access control policy for the web and browsers. This policy has essentially defined a principal model where web sites are mutually distrusting principals [38, 39], and where one site’s script cannot access another site’s content. However, the authenticity of the principal and the integrity of its content are often at question since much of the web is delivered over HTTP rather than HTTPS. Consequently, network attackers can carry out man-in-the-middle attacks and undermine browsers’ access control, even if browsers flawlessly implement the enforcement of the same-origin policy. Such attacks are highly practical today with the prevalence of wireless hotspots and insecurity in the DNS infrastructure [19]. The web requires *end-to-end security* to allow meaningful SOP enforcement in browsers.

HTTPS [31] has the potential to prevent network attacks, but its universal adoption is hindered by its uncacheability at intermediate servers, such as content distribution network (CDN) servers and HTTP proxies, and its performance cost.

Web caching offers significant benefits to web sites and users. It enables web sites to save bandwidth costs and reduce latency for users by outsourcing infrastructure to CDNs and offloading requests to CDN servers. Although CDNs do offer services for HTTPS

content [6], this is at the cost of trusting CDN servers to be man-in-the-middle and losing end-to-end security. Furthermore, such services come with a hefty charge of up to \$3,000 per month plus bandwidth costs [15]. Web cache proxies can also deliver web content significantly faster to large user communities behind gateways or firewalls, such as mobile users. HTTPS content cannot take advantage of these proxies at all today. We observe that much of the web is cacheable (Section 4.1), and we expect significant growth in cacheable web content as rich media proliferates [2]. To achieve an end-to-end secure web, HTTPS is definitely not the complete answer.

In terms of performance, although GMail has recently demonstrated the ability of serving HTTPS content with low overhead using commodity hardware (1% CPU load, less than 10KB of memory per connection and less than 2% network overhead) [26], a general applicability of their solution to other SSL setups is not clear [3]. Due to differences in HTTPS deployments, it might not be trivial for other web sites to replicate Gmail’s performance improvements. Even if the SSL’s server overhead is successfully reduced, it still suffers from lack of in-network caching, thus limiting the performance benefits for the clients.

Fortunately, end-to-end security, cacheability, and performance are not at conflict inherently. End-to-end security encompasses (1) end-to-end authentication (i.e., content comes from the right origin<sup>1</sup>) (2) end-to-end content integrity (i.e., content is not tampered), and (3) end-to-end content confidentiality (i.e., content is kept private). For the browser platform to meaningfully enforce its access control policy, both authentication and integrity are needed, but confidentiality is *not* required. Without confidentiality, the content is cacheable at intermediate web servers. HTTPS provides all three properties simultaneously and is hence not cacheable.

In this paper, we propose *HTTPi* as a protocol to support only end-to-end authentication and content integrity. We advocate that web sites use HTTPS for requests that require end-to-end confidentiality, and HTTPi for all other requests.

This work presents a practical and comprehensive design and implementation of HTTPi that is based on a content-signature-based scheme. While HTTPi requires both browser and server-side modifications, our design does not require changes at intermediate nodes, such as proxies, for caching HTTPi content (Section 2.1). Our design also ensures that progressive content loading in browsers is not hindered by HTTPi, and that this incurs minimal overhead in both computation and bandwidth (Section 2.1). Because signatures can be computed offline and cached for static content, HTTPi has a much lower computational cost compared to HTTPS for web servers.

We further discover that a significant portion of existing HTTPS content can be shared and cached across users (Section 4.1). This

<sup>1</sup>Client authentication is at the discretion of web sites.

indicates that much of existing HTTPS content can be safely turned into HTTPi content to have better performance and the ability of being offloaded to other servers without any loss of security. In fact, many existing HTTPS sites contain HTTP content including scripts and images. Such mixed-content pages often contradict the intent of web sites to defend against network attackers. This is precisely due to the cost of enabling HTTPS for such existing HTTP content. It is much easier to turn HTTP content contained on HTTPS sites into HTTPi content, which will achieve the end-to-end security desired by these sites.

Although we envision a next-generation web with only HTTPi and HTTPS content, HTTP content will undoubtedly exist for a long time. We also provide web developers with an easy way to specify policies of how the three types of content can be safely mixed together (Section 2.2). Furthermore, we observe that the default isolation policy for HTTPi, HTTPS, and HTTP content of the same domain and port does not need to be as strict as the same-origin policy. To this end, we design a new default policy to allow useful interactions across different protocol schemes without sacrificing security (Section 2.3).

End-to-end authentication also requires binding a public key to an origin. Today, such bindings are established through Certificate Authorities. Recent observations have shown weakness in such CA-based binding [16]. DNSSEC can potentially offer a more natural and safer way of binding a domain name to its public key [7]. We will not further discuss this topic in this paper.

We have built an end-to-end prototype to evaluate HTTPi. On the browser side, we implemented the HTTPi protocol for Internet Explorer using IE’s Asynchronous Pluggable Protocol extension mechanism. On the server side, we implemented support for HTTPi requests using an HTTP proxy sitting in front of origin web servers.

Our microbenchmark measurement indicates that HTTPi incurs an acceptable verification and one-time signing overhead, with our unoptimized implementation. This cost is quickly amortized over many requests; for example, a typical web server deployed on Amazon EC2 achieved a 4.06x higher throughput for static content served over HTTPi (and signed offline) than over HTTPS and HTTPi’s throughput is negligibly lower than that of HTTP. To evaluate the efficacy of deploying HTTPi for today’s web sites, we conducted an initial measurement of cacheability of today’s web and found that both HTTP and HTTPS content on today’s web is significantly cacheable. We also present our initial findings on the effectiveness of caching proxies to understand shared caching benefits for web users behind those proxies. Overall, our evaluation suggests that HTTPi is practical to deploy and can offer compelling benefits.

## 2. DESIGN

We set the following goals for the HTTPi design:

- *Guarantee of end-to-end integrity*: Our design ensures that the integrity of the rendered content is always maintained. For example, a network attacker will not be able to inject or remove content, or have adverse impact on browser-side rendering of content.
- *Easy adoption*: HTTPi should be easy to adopt by web sites and should fit seamlessly into the current web infrastructure. In other words, the design should be transparent to the intermediate web servers (such as CDN servers and HTTP web proxies) and should involve minimal changes to the core setup of the servers and the browser.
- *Negligible overhead over HTTP*: The design should incur negligible overhead over HTTP in computation, bandwidth, and user-experienced latency.

Note that there could be scenarios where intermediate servers modify web content, such as for personalization or content filtering in enterprises. Transmitting content over HTTPi instead of HTTP would prevent such modifications. We argue that the guarantee of integrity must be end-to-end, and any intermediate modifications should be explicitly approved by the one of the endpoints (for example, by sharing the private and public key pair of an endpoint).

To guarantee end-to-end integrity and to minimize latency and overhead, we use a content signature-based scheme that allows progressive content loading and at the same time is robust to any injection attacks, as described in Section 2.1. In Section 2.3, we describe the access control policy that browsers should carry out across HTTPS, HTTPi and HTTP content.

For easy adoption, we use the existing HTTP protocol to implement HTTPi so that intermediate web servers can cache HTTPi content seamlessly. Web browsers can show “httpi” in the address bar, but the messages on the wire speak HTTP. We use a new *Integrity* header to indicate the use of HTTPi as the protocol. The integrity header also carries the signature for HTTP headers (excluding the integrity header itself, of course). We use the existing *Strict-Transport-Security* header to prevent stripping attacks (Section 2.1.3) and the existing *X-Content-Security-Policy* header to allow web sites to configure mixed content policies (Section 2.2). Signatures for the HTTP response body are in-band in the body itself. HTTPi’s server-side and client-side implementation is pluggable into the existing setup and uses public interfaces without any need for modifying the core functionality of the server or the browser (Section 3).

### 2.1 Design Overview

A protocol scheme that ensures message integrity needs to satisfy two requirements. First, the identity of the server sending the content needs to be authenticated and second, the content needs to be verified for integrity. HTTPi uses a content signature-based protocol scheme to satisfy these requirements.

In a strawman design, HTTPi could sign the hash of an *entire* HTTP response: The server first creates a cryptographic hash (e.g., SHA1) of the whole response and then signs the hash using the server’s private key. The hash and its signature are then passed to the client along with the response. At the client side, the browser waits for the entire response to arrive, calculates its hash, and compares the value with the signed hash to authenticate the server and verify the response.

A key limitation of this design is that the browser would have to wait for the entire response to arrive before being able to verify the content integrity and dispatch the content for rendering. Consequently, this would disrupt the existing progressive content loading mechanisms in browsers, servers, and the HTTP protocol and the user would experience much longer delay before seeing any content.

We leverage previous work on content integrity [20, 21] to develop our HTTPi design that supports progressive content loading through the use of *HTTPi* segments. While these earlier efforts focused on designing protocol schemes for verification of integrity in streaming systems, our scheme is designed to be used with today’s web applications and browsers. As a result, we solve new problems not addressed in prior work, including compatibility with “chunked” transfer encoding (Section 2.1.1) that is widespread on the web, various content replay attacks, and stripping attacks (Section 2.1.3). Moreover, our work presents a detailed, practical implementation of HTTPi, whereas earlier work focused on theoretical protocol design and offered no implementation details. Before diving into our design for HTTPi, we first provide some background.

### 2.1.1 Existing Progressive Content Loading Mechanisms

Current browsers support progressive loading of web content: as soon as some data arrives from the network, the browser renders it to the user. The amount of data available at a time is determined by the underlying TCP congestion control and the network condition as well as server availability. HTTPS content can also enjoy progressive content loading especially when a stream cipher is selected by web sites.

Complementing browsers’ progressive content loading, servers are also motivated to reduce user wait time and to start sending the response even before completing the processing of a request and therefore, before knowing the entire response body. To this end, servers often use HTTP chunked transfer encoding [18] and encode each piece of available response data into a chunk. A web server typically uses chunked encoding in two scenarios: (1) content is static, however, its retrieval (for example, from the server database) or processing is slow, and (2) content is dynamically generated with a chunk being a logical unit of content for the application. The chunks are sent in separate HTTP responses as soon as they are available. Note that the data of a chunk may not arrive at the client in one shot, but possibly in pieces due to network congestion. Nevertheless, the browser can consume partial chunks progressively.

### 2.1.2 HTTPi Segments for Progressive Content Loading

In HTTPi, the key challenge in supporting progressive content loading is to configure the sensible granularity of content verification. This design should meet the following goals: (1) it leverages browser-side progressive content loading; (2) it is compatible with HTTP chunked transfer encoding; (3) it is resilient to the dynamics of the underlying TCP congestion control, which is unpredictable by servers in an offline fashion; (4) it must allow cacheability; (5) it incurs low overhead.

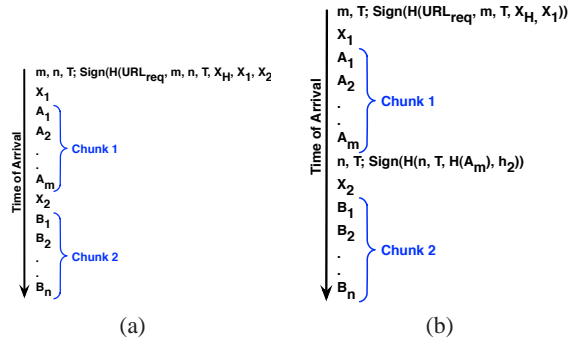
We use *HTTPi segment* to refer to the unit of verification in HTTPi. Let  $S$  denote the size of an HTTPi segment.

Using HTTP chunks as HTTPi segments would still be too coarse-grained. An HTTP chunk can be arbitrarily large and shares the same problem as the strawman solution described above.

A question one may ask is whether a server can predict how much data arrives at its clients. If so, then a server could enable verification for just that data. For a single, live connection, a server can indeed predict data arrivals on the client by obtaining the current TCP congestion control window size and the receiver window size from the network layer. However, because of dynamic network conditions, such prediction would not work well for requests at different times or from different users and would defeat cacheability. In light of this observation,  $S$  needs to be a constant value.

We choose to use the typical TCP segment size (1400 bytes) for  $S$ . TCP segment is the unit of TCP transfer. The rationale here is that the browser will need to wait for *at most* one packet to arrive to receive a full HTTPi segment, perform the verification and render the segment. This wait is as minimal as it can get.

Although HTTPi segment is the unit of verification, it does not need to be the unit of signing. In our design, we amortize the signing cost over multiple segments in the response body. In more detail, whenever a web server has some HTTP response data ready (whether it is the entire HTTP response or an HTTP chunk becoming available), for every  $S$  bytes, we take a hash, then we compute the signature for multiple hashes concatenated in the right sequence. For HTTP headers, we hash each header individually and use a single signature over all hashes. Since browsers do not consume partial header values, we chose not to use the segmented



**Figure 1: Protocol Scheme in HTTPi for (a) static content (b) dynamic content.**  $A_1, A_2, \dots, A_m$  and  $B_1, B_2, \dots, B_n$  represent segments for Chunk 1 and 2, respectively.  $X_1$  and  $X_2$  represent concatenated hashes evaluated over the segments of Chunk 1 and 2, respectively.  $X_H$  represents concatenated hashes over the HTTP headers.  $URL_{req}$  is the requested URL and  $T$  is the time stamp.

design for header fields. We further amortize the signing cost by signing the hashes of HTTP headers along with the hashes of HTTP content using a single signature. We put the signature together with the sequence of the hashes at the beginning of the response body. An alternate way is to put the signature and hashes in an HTTP header. However, our scheme needs to support HTTP’s chunked encoding where chunks after the first chunk do not have header fields. Therefore, we place the signature and hashes inband with the response body.

The decision on when to sign rests with the application and is made based on whether the content being signed is known in advance (i.e., static content), or is generated on the fly (i.e., dynamic content). Figure 1 gives an illustration of our protocol scheme. As can be seen in Figure 1(a), we amortize the cost of signing by using a single signature over segments for all chunks generated for static content (e.g.,  $X_1$  and  $X_2$  in a single signature). Since the content is known in advance, the signature and all corresponding hashes can be pre-computed by the server. For dynamic content, the hashes are computed at the time of content generation. The signature is calculated over all the segments of a single chunk (Figure 1(b)). The sequence of hashes for the headers ( $X_H$ ) is placed only in the first signature. We also place the URL of the requested page ( $URL_{req}$ ) in the first signature and the current time stamp ( $T$ ) in each signature as a preventive measure for certain attacks (Section 2.1.3).

Note that signing can be done in an offline fashion for static content. For dynamic content, this incurs a computation overhead of one SHA1 computation per 1400 bytes, resulting in the bandwidth overhead of just 1.4% (20/1400). The signature overhead is one signature per chunk for dynamic content. We will show in Section 4 that much of the web is static and cacheable and HTTPi incurs negligible overhead over HTTP.

Any segment that fails the integrity check is not rendered. In such cases, we inform the user about the integrity failure and remove the security indicator from the page. For JavaScript, we do not perform progressive content loading because today’s JavaScript engines require an entire script to be received before starting its execution.

### 2.1.3 Security Analysis and Design Enhancements

**Out-of-sequence Segments.** The segment hashes are arranged in a sequence before signing. If a network attacker tries to reorder the segments, it will break the sequence of the hashes and signature verification would fail.



**Injection and Removal Attacks.** Attacker will not be able to launch injection attacks successfully because the injected content will not be verified by the browser. Removal attacks cannot happen to the segment group of a signature for the same reason.

Nevertheless, removal attacks can happen across signature groups (a set of chunks for static content or a single chunk for dynamic content). When HTTP chunks are used by a server, each signature group will have a set of HTTPi segments and a signature for them. A network attacker can remove a signature group without being noticed at the client. To address this issue, we insert the hash of the last segment of the previous chunk at the beginning of the hash sequence of the current chunk (Figure 1(b)); and we insert the header hash at the beginning of the hash sequence of the first chunk.

**Content Replay.** Network attackers could also mix-and-match old content and new content to cause disruptions. Such attacks are prevented in our design by placing time stamp  $T$  in each signature. For HTTPi responses that involve multiple signatures, the browser must verify that the time stamp is the same across all signatures.

The network attackers could alternatively replay a completely different response for requested object. In order to correctly identify the response with the requested object, the client verifies its own value of the requested URL against the signed  $URL_{req}$  value.

**Stripping Attacks.** Both HTTPS and HTTPi are prone to “stripping” attacks that hijack a user’s initial insecure HTTP request and remove redirects to secure content. Although it is possible to notice stripping attacks by manually checking the browser security indicators, users often ignore these indicators [34]. The HTTP Strict Transport Security protocol (HSTS) prevents these attacks by allowing web sites to specify a minimum level of security expected for connections to a given server. The policy can be delivered via HTTP header [23]. To prevent attacks on the user’s very first visit to the site, the policy can also be delivered via DNSSEC [25]. We use an extension to HSTS, `allowHTTPi`, to allow servers to specify HTTPi as the minimum level of security. The `allowHTTPi` token is appended to the server’s existing `Strict-Transport-Security` policy declaration. Older browsers that do not support HSTS will ignore this header, while older browsers that support HSTS but not our extension will default to HTTPS for all content.

**Denial of Service.** HTTPi is limited in its capability to handle denial of service attacks, where a network attacker strips off the integrity header from the response that requires integrity as specified by the application (Section 2.2). As a result, the content would not be rendered by the browser. Additionally, the attacker can allow some segments to be rendered, while preventing subsequent segments to arrive through to the browser. This could potentially corrupt the internal logic of the application. For example, the attacker can strip off JavaScript that changes the content of the page and as a result, the page remains rendered in its original form. One possible countermeasure to this attack is to use a time out for inter-segment arrival at the client and raise an integrity failure alert after the expiration of the timer. However, it would require an estimation of the typical inter-arrival time for each client, which might not always be accurate. In our design, we allow the browser to wait infinitely for the packets to arrive. If the user clicks on stop, we alert the user that the content is not complete. Since we do not execute JavaScript till it is fully received, partially rendered JavaScript would not be an issue for the integrity of the site.

## 2.2 Mixed Content

The mixed content condition occurs when a web developer references an insecure (HTTP) resource within a secure (HTTPS) page. Such references create vulnerabilities that put the privacy and integrity of the otherwise secure page at risk, because the insecure

content could be modified in network transit. Scripts are particularly problematic because they acquire the principal origin of the including page, allowing malicious scripts to read or alter the content that was delivered over the secure connection. These types of vulnerabilities are becoming increasingly dangerous as more users browse untrusted networks and attackers improve upon DNS poisoning techniques and weaponize exploits against insecure traffic.

Browsers differ in their mixed content handling. Internet Explorer prompts the user before displaying mixed content, while Firefox and Google Chrome show a modified browser lock icon. From a security standpoint, the best behavior would be to block all insecure content in secure pages without prompting the user. The latest beta release of IE9 enforces this behavior on scripts and stylesheets, but not images; this policy is similar to the one proposed by Gazelle [39]. However, this option of automatically blocking insecure content has some serious compatibility implications. It might potentially confuse the user, since pages that rely on insecure resources could appear broken. In the worst case, the user might think the broken pages indicate a bug in the browser and subsequently switch to an older version of the browser or to a completely different browser to get unbroken pages.

We argue that mixed content vulnerabilities should be fixed by the web developers, both for security and user-experience reasons. The web developers have a better understanding of the impact that embedded content can have on the security of their site. Additionally, they are in much better position to develop a user-friendly fallback mechanism for their site in case some content fails security checks and hence is not rendered.

By default, we require that all active content embedded in HTTPi and HTTPS pages, such as scripts and stylesheets, be rendered over HTTPi or HTTPS. To allow web applications to customize this default behavior, we use an HTTP header that is compatible with Content Security Policy (CSP) [36] header to specify the server’s end-to-end integrity requirements for dependent resources. The CSP policy syntax is convenient for our purposes as it already allows sites to specify which origins they want to include content from.

An example policy is as follows:

```
X-Content-Security-Policy:  
  allow https://login.live.com  
  httpi://*.live.com:443
```

The above example informs the browser that all embedded objects from `login.live.com` should be retrieved over HTTPS and content from all other subdomains of `live.com` needs to be downloaded over HTTPi. If the servers hosting the embedded content do not support the corresponding protocol, then the content is considered unsafe as per the web page’s requirements and hence should not be rendered by the browser. Our design also supports specification of integrity requirements at a finer level, i.e., at the level of object types or specific objects themselves. However, the web application should be careful in specifying such finer policies as it increases bookkeeping at the server. It also has the potential to break existing interactions within the embedded content if the policies are not correctly specified.

The CSP syntax provides an ideal mechanism for the web developers to handle mixed content. It does not require web applications to change their code by explicitly modifying all insecure references of embedded objects. Even if web developers decide to modify their code, it might not be sufficient. A secure (HTTPS or HTTPi) URL can still return a redirect to an insecure resource, which could be difficult to determine by examining the DOM alone. Additionally, a script delivered over a secure channel could still make references to insecure content. In our design for HTTPi, the browser enforces the policies specified by CSP for all statically or dynami-

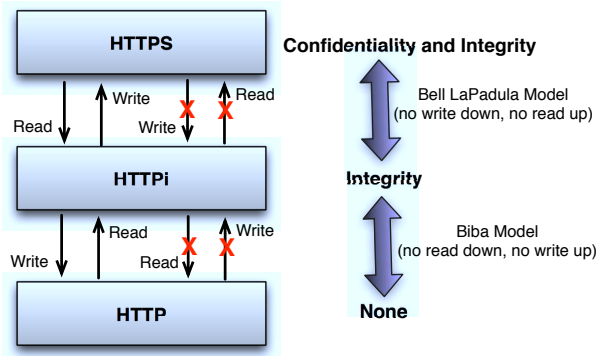


Figure 2: Interactions in Mixed Content Rendering.

cally generated URLs.

### 2.3 Access control across HTTPS, HTTPi, and HTTP content

HTTPi content can be embedded in an iframe through the use of the “httpi” scheme, such as `<iframe src=“httpi://a.com/”>`, or through the use of an additional iframe “integrity” attribute, such as `<iframe src=“http://a.com/” integrity>`. The former has the consistent presentation with other protocol schemes. The latter has the benefit of backward compatibility; on an older browser, HTTPi content would simply render as HTTP content. Note that no matter what the representation is, the network messages still speak HTTP to be backward compatible with the existing web caching infrastructure.

The Same Origin Policy labels the principals with the origin defined as the triple of `<protocol, domain, port>` [38, 39]. Therefore, content from the same domain and port number but with different protocol schemes is rendered as separate principals. They can only communicate explicitly through messages (i.e., `postMessage`) [8].

In this subsection, we consider the default interaction and access control model for HTTPS, HTTPi, and HTTP content served from the *same* domain and port. For example, a top-level HTTPi page may embed two iframes, one containing HTTP content and the other containing HTTPS content; and all three pages are from the same domain and port. While following the SOP is safe for such scenarios, it disallows all interaction among HTTP, HTTPi, and HTTPS content. Rather than accessing the DOM objects directly, developers would be forced to redesign such interaction with asynchronous `postMessage`-based protocols, which may be hard to design correctly, as illustrated by recent flaws found in Facebook Connect and Google Friend Connect [22]. As a result, a developer may be discouraged from converting some content on an HTTPS site into HTTPi to benefit from its cacheability properties.

As a concrete example, consider an online shopping site that is rendered over HTTPS to protect users’ private data such as credit card information. The site presents users with a map to select a site-to-store pick-up location during checkout. It may be desirable to deliver the store information and map content over HTTPi, but this raises a problem of allowing the HTTPS part of the site to read the store selection made by the user, an interaction that would be disallowed by SOP. As a result, the site’s developers may be forced to refactor their code to use `postMessage`.

We observe that the SOP semantics are more restrictive than actually required to ensure security for such scenarios. Our goal is to allow legitimate communication while preserving the security semantics, namely the confidentiality and/or integrity, of the ren-

dered data. Our default communication policies are inspired by the combination of the Bell LaPadula [10, 11] and Biba [12] models. It is important to note that our goal is *not* to enforce information flow invariants often associated with those models (e.g., frames of any origin can already freely communicate via `postMessage`), but rather to use these models to determine a secure and convenient *default* isolation policy for our setting. We summarize these models as the following set of rules:

#### Bell LaPadula model (for confidentiality):

- The Simple Security Property: a subject at a given security level may not read an object at a higher security level (no read-up).
- The \*(star) property: a subject at a given security level must not write to any object at a lower security level (no write-down).

#### Biba model (for integrity):

- The Simple Integrity Axiom states that a subject at a given level of integrity may not read an object at a lower integrity level (no read-down).
- The \*(star) Integrity Axiom states that a subject at a given level of integrity must not write to any object at a higher level of integrity (no write-up).

In view of these models, we represent the three protocols (HTTP, HTTPS and HTTPi) by two confidentiality levels ( $C_{high}$  and  $C_{low}$ ) and two integrity levels ( $I_{high}$  and  $I_{low}$ ), which models the high and low requirements for confidentiality and integrity, respectively. HTTPS can be realized by the tuple  $\langle C_{high}, I_{high} \rangle$ , HTTPi by  $\langle C_{low}, I_{high} \rangle$  and HTTP by  $\langle C_{low}, I_{low} \rangle$ . Using this model, we define the access control rules across HTTP, HTTPi, and HTTPS as follows:

**HTTPS and HTTP.** HTTPS’ confidentiality label  $C_{high}$  is higher than HTTP’s confidentiality level  $C_{low}$ , thus resulting in “no read up, no write down” requirement of the Bell LaPadula model. The integrity levels of HTTPS and HTTP,  $I_{high}$  and  $I_{low}$  respectively, with  $I_{high} > I_{low}$ , results in “no write up, no read down” condition of the Biba model. Combining these two requirements results in no reads or writes to either side being allowed between HTTPS and HTTP. This derivation is consistent with the SOP.

**HTTPi and HTTP.** Since confidentiality levels of HTTPi and HTTP are equal, only the integrity levels enforce the “no write up, no read down” policy from the HTTPi content to HTTP resources (Figure 2). Firstly, this means that a script belonging to the HTTPi principal can write to the HTTP part of the page without reading its content. One reason to prevent an HTTPi script from reading HTTP content is to prevent the HTTP input from influencing the logic within the HTTPi content. However, an HTTPi script might still desire to read the HTTP page to identify the DOM element to write to. So, our requirement is to allow the read operation on the HTTP content without allowing the logic of HTTPi content from being affected. One way to realize this is by performing complete information flow check in the HTTPi code, which might not be practical. We use an alternative approach in which the HTTPi content itself writes the code for reading the HTTP content, and this code is injected into the HTTP content. This injected code runs within the HTTP principal and hence can freely read and write to the content. Since HTTPi relinquishes the transferred code to the HTTP integrity level ( $I_{low}$ ), that code cannot affect the logic of HTTPi’s own code, though it still can read from HTTPi content. Secondly, HTTP can read the HTTPi content, but cannot write to it. We realize this in our design by providing only a shadow copy of the HTTPi content to HTTP, with no direct reference to real HTTPi objects.

**HTTPS and HTTPi.** Since HTTPS and HTTPi integrity lev-

els are equal, only the confidentiality levels force the “no read up, no write down” rule from HTTPS to HTTPi resources (Figure 2). Both read and write operations can be realized similarly to the previous scenario. We allow HTTPi content to write to HTTPS since the code for HTTPi is at the same integrity level as HTTPS content and written by the same developer (since they have the same domain). HTTPi scripts can write the code for reading the HTTPS content into the HTTPS’ DOM and effectively, that code becomes part of the HTTPS principal. This allows reading of the HTTPS code by the injected code without leaking any of the read data back to HTTPi’s main code. For reading HTTPi content without allowing any write, a shadow of the HTTPi’s DOM is provided to the HTTPS. Coming back to the shopping site example earlier in this section, this rule would allow HTTPS content to read the store selection made by the user and correspondingly send the merchandise to the selected store.

### 3. IMPLEMENTATION

HTTPi requires both the client browser and the hosting server to adhere to the protocol. Accordingly, our implementation consists of server-side and client-side modules. Figure 3 shows the high-level architecture of our system. Our server-side implementation consists of an HTTPi Transformer, which implements all HTTPi-related interactions on the server side, including content hashing, segmentation, and a handler for appending integrity policy requirements to HTTP responses. Our client-side implementation centers around three modules that we add to Internet Explorer 8: (1) an HTML content filter that transforms a given page to adhere to integrity policy requirements, (2) an HTTPi protocol that handles the client-side processing of HTTPi content, and (3) a module that provides JavaScript and DOM interposition to enforce our mixed-content access control policies. In this section, we describe each of these modules and the associated implementation challenges in turn. Overall, our implementation consists of 1,100 lines of server-side code, and 3,500 lines of client-side code.

#### 3.1 Server-side Implementation

We explored two options for implementing the server-side component of HTTPi, with the options differing in their deployment tradeoffs. First, we extended the IIS 7 web server with a C# module for HTTPi, called HTTPi Transformer, that encapsulates the functionality to generate HTTP responses with signatures and content hashes that adhere to the HTTPi protocol. Although we chose IIS, similar module functionality is available for other web servers. This option is useful if the server is willing to immediately integrate HTTPi functionality into their current setup. It also has obvious performance benefits as the module is closely coupled with the functionality of the web server.

In our second deployment option, we integrated the HTTPi Transformer into a web proxy that translates typical HTTP responses into HTTPi responses by embedding all the hashes and signatures needed by HTTPi. We leveraged the public-domain Fiddler web debugging proxy [27] and its FiddlerCore [17] extensibility interfaces. This option is independent of web server implementation and allows servers to continue supporting HTTP as the delivery protocol for backward compatibility, while switching to the HTTPi protocol for requests that pass through the proxy. It eases deployment, since the proxy can be deployed anywhere in the network and guarantees integrity between the proxy and a compatible browser. This could be desirable for corporations that do not require integrity checks for intranet users, while ensuring integrity of their sites for external users.

For our evaluation, we used the latter option of having a network

proxy, because (1) it allowed us to test our prototype against publicly deployed web sites without having any control of their web servers, and (2) it allowed fair comparison of HTTPi with HTTPS and HTTP (Section 4.2.3) by cleanly switching to a desired protocol between the client and the proxy even when the backend server did not support the protocol.

### 3.2 Client-side Implementation

#### 3.2.1 Filtering content to enable HTTPi

We expect that origin servers would generate new content with the right “httpi” URIs for the content that requires integrity. In any case, our design ensures that mixed content policies are enforced by verifying the URIs against the policies. Instead of requiring the servers to change the URIs in their existing content, our implementation of HTTPi performs the required filtering to enforce the mixed content policies. The HTML content filter module is invoked for every HTML response received at the browser that is associated with a Strict Transport Security policy or Content Security policy. This module modifies HTML content to ensure that it adheres to the minimum security levels specified in STS and CSP. For example, all object links on a page are transformed to corresponding HTTPi links by modifying the protocol field in the URL. Since the HTML content filter is invoked before the page is rendered in the browser, this design allows the HTTPi protocol handler to be associated with all such links and hence ensures that the HTTPi handler is invoked when the browser requests those links during rendering. We implemented this module by using IE’s public MIME filter COM interfaces [1] and subsequently registered it as a filter for HTML content.

One limitation of this approach is that it may miss dynamically-generated links where the URL is constructed by JavaScript at runtime. We are currently working on solving this by performing HTTPi redirection to the time of actual HTTP requests; our evaluation is independent of this implementation enhancement and was performed without it.

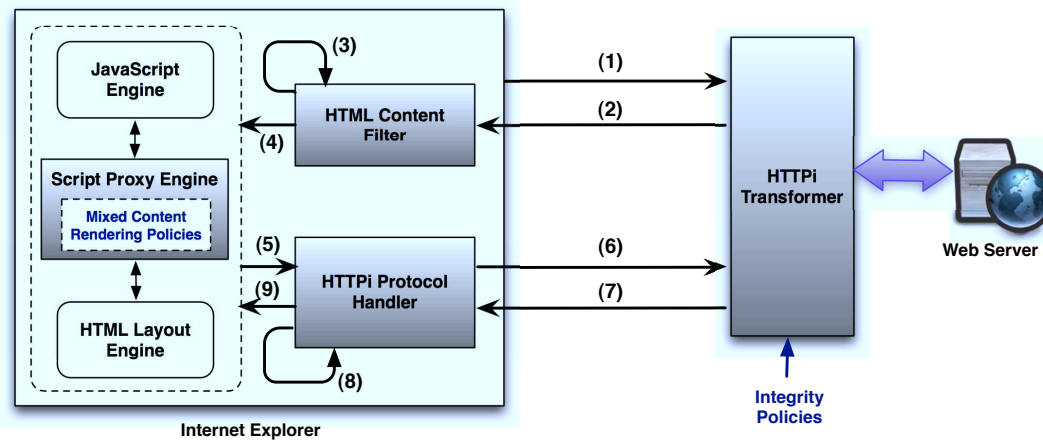
#### 3.2.2 HTTPi Protocol

The HTTPi protocol handler encapsulates all client-side handling of HTTPi content and is automatically invoked by the browser when an HTTPi link is encountered by the browser’s rendering engine. Upon invocation, it makes an independent HTTP call to the server to retrieve the content. It then verifies the integrity of the content in segments using the mechanism described in Section 2. Once the integrity of a particular segment is verified, its content is released to the browser’s rendering engine for progressive loading.

We implemented this module as an asynchronous pluggable protocol (APP) [1] IE module associated with the HTTPi protocol. Even though IE provides this generic protocol extension point, implementing a general-purpose protocol with minimal performance overhead is challenging. IE’s internal logic is well-optimized for HTTP and HTTPS, which makes a comparably performant web protocol difficult to implement. A considerable time and effort was spent on making our code as optimal as possible by parallelizing various operations such as network read and signature verification. Despite our limited knowledge of IE’s internal optimizations and with the handicap of using a generic interface, we were still able to achieve acceptable performance as compared to HTTPS and HTTP (Section 4.2.3).

#### 3.2.3 Access control for mixed content

Another big challenge for our implementation was to customize SOP to include our mixed-content access control policies. Unfor-



**Figure 3: High-Level Architecture of our HTTPi Implementation with the operational steps to retrieve content over HTTPi as follows:** (1) IE makes an initial request for a specific page. (2) Server-side proxy identifies that the request is for a HTTPi-enabled resource and appends integrity policy headers to the response. (3) HTML content filter processes the response by modifying URLs that match STS policies to point to their corresponding HTTPi links. (4) HTML content filter releases the modified response to IE’s rendering engine. (5) The HTTPi protocol handler is invoked for every HTTPi object encountered during rendering. (6) The HTTPi protocol handler makes a HTTP call to the server requesting the object. (7) Server-side proxy traps the request, makes an independent HTTP call to the backend web server to get the response, hashes and signs the response, and returns it back to the HTTPi protocol handler. (8) The HTTPi protocol handler verifies the signature and hashes corresponding to the different segments in the response. (9) Successfully verified segments are passed to the rendering engine for progressive loading. The Script Engine Proxy (SEP) subsequently mediates all mixed-content interactions while a web page renders.

fortunately, IE does not allow changing the code for SOP with public APIs. As a result, the only alternative was to implement our solution as an additional layer on top of the existing SOP and then find a way to enforce mixed-content policies within the limits imposed by the existing SOP logic. This certainly made our implementation more difficult.

To solve this problem, we use a two step approach. In the first step, we modify the security origin (origin is defined as the tuple `<protocol, domain, port>`) of all objects on the web page by changing the protocol field to HTTP, i.e., the one with the lowest integrity and confidentiality level. This is achieved by providing a custom implementation for the `IInternetProtocolInfo` interface [4] from within the APP for HTTPi. Note that changing the security origin of an element does not affect the URL associated with that element.

As per the SOP, all the objects on the page can now interact without restriction. Our second step is to enforce access control rules or policies that govern such interactions. We build on our earlier work [35, 38] that implements a JavaScript engine proxy (called script engine proxy or SEP): SEP is installed between IE’s rendering and script engines, and it mediates and customizes DOM object interactions. SEP is implemented as a COM object and is installed into IE by modifying IE’s JavaScript engine ID in the Windows registry. We extend SEP to trap into all invocations (read or write) across the page’s objects and ensure that our mixed-content access control policies (Section 2.3) are enforced. We use the URLs associated with the accessing object and the object being accessed in making our access control decision. The two-step logic that governs the access control enforcement in our implementation can be summarized as follows:

- If the original origins of the caller and the callee objects differ in `domain` and/or `port`, then the browser would prevent any interactions across them in accordance with the SOP.
- If the original origins of the caller and the callee objects dif-

fers in only `protocol`, the SOP would allow the objects to interact (as we modify the protocol of the security origin to HTTP). In this case, we mediate the interaction within our customized SEP to enforce our access control policies.

The read operation is straightforward: SEP allows the caller to have read access to the callee’s objects. The write operation could be implemented in a similar fashion; however, some writes must first access an object to which the write subsequently occurs. For example, if the caller wants to write content to a specific element on a callee object, it might need to read the handle to that element using functions such as `getElementById` or `getElementsByName`. However, if the caller only has write privileges with no read access, it cannot make such calls and hence cannot know where to write the content.

We solve this problem by introducing a new JavaScript function `writeUsingCode`, which is interpreted by our SEP implementation; the browser’s JavaScript engine does not need to understand this function. Instead of directly making read calls looking for an element of the callee object, the caller uses the function to pass the JavaScript code that encapsulates such read calls and the subsequent write call to the corresponding element. The SEP intercepts this function call and makes calls to the underlying JavaScript engine to execute the code with the origin of the callee object. Any unintended feedback mechanism introduced by this code is prevented by SEP’s access control policies.

## 4. EVALUATION

We have implemented a HTTPi system that works end-to-end. We used our proxy-based implementation as a server-side HTTPi endpoint to verify our system for correctness against a number of popular web sites, such as Google, Bing Maps, and Wikipedia. In each case, the browser successfully rendered the web pages and all integrity checks were correctly included at the server and verified at the browser. Any tampering of the web page in the network was



Protocol	Total Objects		Publicly Cacheable Objects	
	Count	Size	Count	Size
HTTP	346,629	1532 MB	251,826 (72.65%)	1385 MB (90.41%)
HTTPS	5,036	21.95 MB	3,659 (72.66%)	19.39 MB (88.33%)

**Table 1: Measurement of publicly cacheable web content from the top 1000 Alexa sites.**

correctly detected and failed the integrity check at the browser. We evaluated the access control interactions for mixed content by developing a set of custom web pages that included such interactions. Our system correctly enforced the access control policies for such interactions.

Next, we provide experimental evidence to support our claim that today’s web sites can benefit from cacheability enabled by HTTPi. To this end, we first perform a web cacheability study to answer two questions: (1) what web sites have cacheable content, and (2) what users are taking advantage of shared caches on the web. Next, we evaluate the performance of our prototype by micro-benchmarking its operations and by comparing its overhead to that of HTTPS and HTTP.

## 4.1 Study of Web Cacheability

With HTTPi, web sites decide what content uses HTTPi as the underlying mechanism of transport. Therefore, any content that web sites currently allow to be cached by intermediate web servers, such as CDNs and web caches, becomes an ideal target for HTTPi. To better estimate the amount of web content that could benefit from the use of HTTPi, we performed a cacheability analysis on the top 1,000 Alexa sites that includes both top-level pages and embedded content on the sites visited. We analyze the HTTP caching headers, such as `Cache-control`, `Expires`, `Pragma`, etc., to decide what content is deemed cacheable according to the HTTP specification [18].

**Experimental Setup.** To facilitate automatic analysis for a large number of URLs, we used a customized crawler from our earlier work [35], which utilizes IE’s extensibility interfaces to completely automate the browser’s navigation. To invoke functionality beyond a site’s home page, the crawler uses simple heuristics that simulate some user interaction, such as clicking of links and searching form submissions. We restrict all simulated navigations to stay within the same origin as a site’s home page. We monitor the browser’s network traffic in a proxy to intercept all HTTP/HTTPS requests and analyze HTTP headers relevant to web caching. The proxy is included as a trusted certificate authority at the browser in order to allow it to intercept the HTTPS traffic and inspect its content [27].

**Prevalence of cacheable content.** Table 1 shows the results of our web cacheability experiment. Note that our results only consider content that is marked as public and excludes any private content that is user-specific and hence is intended to be cached only at the user’s browser. As we can observe from the table, a large majority of the web content is rendered over HTTP with more than 98% of the objects that we observed being HTTP objects. We found that approximately 73% of these objects are cacheable. The cacheability is higher in terms of content size, with more than 90% of total HTTP content size (of all objects) being cacheable, indicating that the web applications typically want larger-sized content, such as images, to be cached in the network. The limited number of HTTPS objects that we encountered follow a similar trend with a large number (73%) being cacheable objects. The presence of a considerable number of public, cacheable HTTPS objects is an indication that web applications intend to cache objects in the web, but are discouraged by the lack of security provided by HTTP. They

are left with no choice but to trust the CDNs for this type of content. If only integrity of the content is desired, HTTPi presents itself as an ideal alternative for these HTTPS objects.

**Presence of in-network caches.** To see how many users are benefiting from web caches today, we measured the prevalence of forward caching proxy servers, which are a significant source of in-network caching. More specifically, we conducted an experiment to determine how the country and the user agent affects whether a forward network proxy is being used. We used rich media web ads as a delivery mechanism for our measurement code, using the same ad network and technique previously demonstrated in [24]. We spent \$80 to purchase 115,031 impressions spread across 194 countries. Our advertisement detected forward proxies using XML-HttpRequest to bypass the browser cache and store content in the network cache. Overall, 3% of web users who viewed our ad were using a caching network proxy. However, some countries had a significantly higher fraction of users behind network proxies. Popular countries for forward proxies included Kuwait (63% of 372 impressions), United Arab Emirates (61% of 624 impressions), Argentina (11% of 1,875 impressions), and Saudi Arabia (10% of 4,248 impressions). We also observed higher usage of forward proxy caches (11%) among mobile users, although these users accounted for only 0.1% of the total impressions in our experiment.

**Relevance to HTTPi.** Our results demonstrate that cache proxies are still prevalent and useful today, particularly for large user communities, such as a whole country of people behind a single firewall and mobile users behind cellular gateways. HTTPi can take advantage of these proxies while offering end-to-end security at the same time.

## 4.2 Performance Evaluation of HTTPi


We evaluate the performance of HTTPi in two steps. First, we perform micro-benchmarking of various stages of the protocol and analyze the parameters that determine HTTPi’s performance. Second, we analyze the end-to-end performance overhead of HTTPi over existing HTTP and HTTPS protocols.

### 4.2.1 Experimental Overview

Ideally, we would run performance experiments on real web sites deployed on the web. However, current web servers do not understand the HTTPi protocol, and many servers host an HTTP version of a site but not HTTPS. To overcome this, we used our modified server-side Fiddler [27] proxy (Section 3.1) for proxying all requests from the client to the backend server, and converting HTTP requests from the origin server into HTTPi or HTTPS requests to the client, as necessary for our experiments. This setup allows us to measure the cost of using HTTPS and HTTPi for web pages that are currently hosted over HTTP.

We use the end-to-end response time as the measurement criterion, defined as time between the instance at which a URL is submitted at the browser and the instance at which the corresponding page is fully rendered. To remove any discrepancies that might arise from fetching content from the backend server due to inconsistent network conditions, we deduct the data fetching time at Fiddler from the total end-to-end response time. This gives us an es-





**Figure 4: Micro-benchmarking various operations in HTTPi for a 836KB web page, using 512Kbps network bandwidth.**

timate of the end-to-end response time with Fiddler acting as the server. For a fair comparison, we also perform similar deductions for HTTP and HTTPS.

For our experiments, we use SSL certificate size of 1024 bits. Even though there is a push on the Internet to move towards 2048-bit certificates, many of the popular sites such as Gmail still use 1024-bit keys. Additionally, it makes HTTPi’s performance estimates to be conservative in comparison to HTTPS, as HTTPS will perform worse for 2048-bit keys.

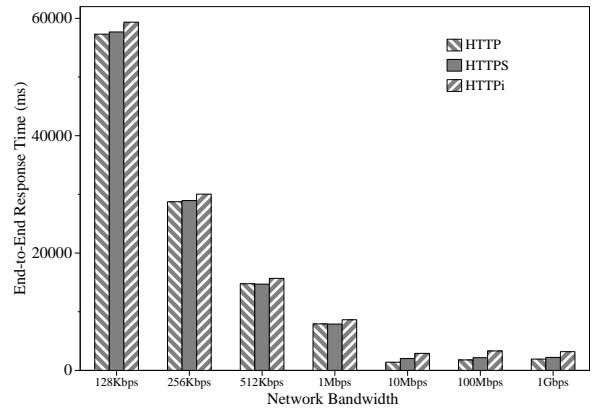
Using the Akma network delay simulator v0.9.129 [5], we simulated various network conditions to understand their performance impact on end-to-end response time. We simulate the incoming and outgoing connections to have equal bandwidth and fixed their queue sizes at 20 packets. We run our delay simulator on the server side to cap the server throughput to a desired bandwidth. We deploy our server-side Fiddler code on a Windows 7 machine, with an Intel 2.67 GHz Core i7 CPU and 6 GB of RAM. The client runs on a Windows 7 machine, with an Intel 2.4GHz quad-core CPU and 4GB of RAM. All experimental results are averaged over 10 trial runs.

#### 4.2.2 Micro-benchmarks

To understand sources of overheads in our system, we instrumented our HTTPi implementation to measure latencies of various operations, and used a simulated network bandwidth of 512Kbps to load an 836KB HTML page in our HTTPi-enabled browser, with the size picked to maximize measurable overhead and to observe effects of HTTPi’s segmentation. Figure 4 breaks down the delays contributing to the end-to-end response time, which we measured to be 15.7 sec.

We find that a large fraction of the total time is spent reading content from the network (bar 7 in Figure 4), which is an expected behavior for slower networks. The overhead costs of hashing all content segments (bar 2) and signing these hashes with a 1024-bit key (bar 3) on the server side is very small. Here, the RSA signature is calculated on a fixed-size single SHA1 hash of 20 bytes (Section 2); this takes just 3ms. Since the header value sizes are much smaller as compared to the content body, both the time to set the header integrity content (hashing and signing) on the server (bar 1) and time to verify it on the client side (bar 4) is low.<sup>2</sup> On

<sup>2</sup>Note that we do not perform any segmentation for headers. For our measurements, we specify two headers, `Server` and `Content-Type`, to require integrity. This time cost will vary



**Figure 5: End-to-end response time as a function of the network bandwidth available to the client, measured for a 836KB page. Note that these results do not include performance benefits due to caching for HTTP and HTTPi.**

the client side, the signature verification time (215 ms, bar 6) is a more significant source of overhead. It is considerably higher than the cumulative hash verification time for all content segments (51 ms, bar 5), supporting our design of using a single signature over multiple segment hashes. The time to pass data from our client-side HTTPi protocol handler into the browser’s rendering engine (bar 8) is also considerable; although it is not specific to HTTPi and would also be incurred by other protocol handlers in the browser, native protocols like HTTP are more optimized in our browser for this step, as we discussed in Section 3.2.2.

In summary, we find that the major HTTPi components (bars 1-6) constitute only 295 ms (1.8%) of the end-to-end response time for this microbenchmark, with largest overhead coming from client-side signature verification.

#### 4.2.3 Comparing HTTPi to HTTP and HTTPS

In this section, we compare HTTPi’s performance to that of HTTP and HTTPS and answer two questions: (1) Is the user-perceived latency acceptable for the data received over HTTPi, and (2) What is the performance impact of running HTTPi and the hashing and signing load it incurs on a web server?

**User-perceived latency.** We compared the end-to-end response time for our 836KB test page rendered over HTTPi, HTTPS and HTTP. Figure 5 shows the results of our experiments performed over different network bandwidth conditions. Note that the performance results do not include caching, and only evaluates the first of potentially many requests for this page. Evaluating performance of a particular cache is not a goal of our experiments and has been previously well studied [40, 41]. We see that HTTPi incurs minimal overhead over both HTTP and HTTPS, and this overhead is consistently within 0.7-2.0 seconds over both HTTP and HTTPS for different network bandwidths. Since this value does not vary much with network bandwidth, we believe our implementation is successful in approximately matching the network optimizations of HTTPS and HTTP. We believe that there is still ample room for client-side optimization as we discussed earlier in Section 3, and this will certainly reduce the total overhead of HTTPi (since client-side overhead is not negligible as shown by our micro-benchmarking experiments).

**Web Server Throughput.** Our server throughput measurements are performed using httpperf [28], an HTTP performance measure- according to the number of headers for which integrity checksum is set.

Experiment	HTTP	HTTPi	HTTPS
Bare-metal Setup	3320	3318	2503
Amazon EC2 Setup	2757	2732	678

**Table 2: Impact of HTTPi and HTTPS on server throughput in responses/sec.**

ment tool. The experiments are performed using two different setups that closely represent typical real-world web deployments:

- Our first setup consists of an IIS server that is hosted on a bare-metal Windows 7 machine, with Intel 2.67 GHz Core i7 CPU and 6 GB of RAM. The Linux client machine running `httperf` is connected to the server by a 1Gbps network with negligible latency.
- Our second setup is cloud-based; we use a virtual Windows 2008 Server image on Amazon EC2. At the time, this image was the only publicly available image that came pre-installed with IIS 7. It is a “high-CPU medium” instance with 5 EC2 compute units with 1.7 GB of RAM (the fastest instance that was available for this image). This setup mimics a typical EC2 user who wants to host a web server. `httperf` is executed from a Linux EC2 instance in the same region, using EC2-private LAN with negligible network latency.

We use an experimental HTML page of size 4.8 KB, which represents a typical size of a page with no embedded links. We arrived at this page size based on the web estimates that put total page size at 170KB (median) and number of objects per page at 37 (median) [29]. For each page, we increased the offered load on the server until the number of sustained sessions peaked. We found that the server was CPU-bound in all cases. Each session simulated one request to the web page.

Table 2 shows a summary of our results. HTTPi incurs negligible degradation (less than 1%) of throughput compared to the original HTTP page. In comparison, the throughput drop was substantial when using HTTPS, with our bare-metal experiment reporting 25% and EC2 experiment showing 75% drop in the throughput. This drop is attributed to the heavy CPU load for the SSL handshake. Our bare-metal experiment shows a lesser drop since it has a considerably faster CPU, which handles the load better. Overall, these results demonstrate that web servers can have a significant performance incentive to use HTTPi instead of HTTPS.

## 5. RELATED WORK

Prior work has explored a number of integrity protection techniques. A proposal on authentication-only ciphersuites for PSK-TLS [13] describes a transport layer security scheme for authentication and integrity, with no confidentiality guarantees. However, this proposal requires a shared secret between each client and the server to key the hash, making it impractical to share the key with all the clients of the application. SHTTP [32] is another proposal for a content-signature-based protocol that unsuccessfully competed with SSL and HTTPS. Our work builds on SHTTP’s signature mode of operation and develops a practical and comprehensive solution by additionally addressing progressive content loading, mixed content handling and the associated security.

Web tripwires [30] verify the integrity of a page by matching it against a known good representation of the page (either a checksum or an encoded full copy of the page’s HTML). It uses client-side JavaScript code to detect in-flight modifications to a web page. However, web tripwires have high network overhead (approximately 17% of the page size), which could hinder the end-to-end response time, especially for slower networks. Moreover, web tripwires can be identified and disabled by an adversary, and they cannot detect

full-page substitutions. In contrast, HTTPi is cryptographically secure and can prevent any type of integrity breaches. HTTPi also has a much lower network overhead cost as compared to web tripwire. Finally, web tripwires focus on *detection*, while HTTPi focuses on both *detection* and *prevention*.

Other research has proposed cryptographic schemes for web content integrity [9, 20]. While we share some commonality with this work in integrity computation, our system differs in three significant ways. First, our design is more robust against attacks like stripping attacks and content replay. Second, we design HTTPi to be practical for today’s web and address problems such as mixed content treatment, compatibility with “chunked” transfer encoding, and access control across HTTP/HTTPi/HTTPS content, none of which are considered in prior work. Third, we go beyond protocol design and also offer a full practical implementation and evaluation of HTTPi for a real-world browser, while earlier research lacks any implementation details.

Stubblefield et al. [37] proposed mechanisms to improve SSL’s performance. While their WISPr system shares HTTPi’s motivation of supporting in-network caching while preserving integrity, it is designed for another content delivery protocol (subscription-based), rather than for use in existing web sites. WISPr constructs an HTTP page that embeds the encrypted version of the original page; this page can be cached in the network. However, a client needs to download a key from the server in order to decrypt the content, and WISPr only works for static content. In contrast, HTTPi is readily compatible with existing web sites, it supports static and dynamic content, and it adds support for progressive loading and mixed-content scenarios common on the web. Whereas no evaluation details are provided for WISPr implementation, we showed that HTTPi is practical in Section 4.

HTTP provides a Content-MD5 header [18] that can carry the MD5 signature of the complete page. This header could be useful in providing basic page integrity, but suffers from many weaknesses if used by itself. For example, a network attacker can modify the header since it is not authenticated, and the attacker can completely drop the header without the client knowing about it. In contrast, HTTPi provides authentication by signing content hashes, and since it specifies the requirements for a page using HSTS in advance, the client can easily detect whether the required integrity content is dropped by network attackers. Additionally, with HTTPi, integrity is evaluated over smaller-sized segments, which has performance benefits (see Section 2.1) over the entire-page approach used in the Content-MD5 header.

The YURL [14] specification defines an alternative server identification and authentication mechanism that does not depend on centralized authorities like the DNS or PKI. A YURL identifies a site using the site’s public key fingerprint and the web site owner owns the CA fingerprint. However, like HTTPS, and unlike HTTPi, the proposed YURL-based protocol *httpsy* [14] precludes content from being cached at web proxies.

## 6. CONCLUSIONS

We envision HTTPi to complement HTTPS to bring end-to-end security to the entire web. Only when there is end-to-end security, the browser platform and the web are able to have a collectively secure overall system.

We advocate the part of web that does not have end-to-end security today to adopt HTTPi which incurs negligible performance overhead over HTTP and enjoys the benefit of CDNs and cache proxies just as HTTP. For existing HTTPS content, our study indicates that its significant portion is cacheable and can also gain significant performance and caching benefit from employing HTTPi.

## 7. ACKNOWLEDGMENTS

We would like to thank Shai Herzog and Gil Shklarski for their insightful discussions and support. We are also grateful to Carl Edlund, Jim Fox, Justin Rogers, and Ali Alvi for their help with IE instrumentation.

## 8. REFERENCES

- [1] About Asynchronous Pluggable Protocols. [http://msdn.microsoft.com/en-us/library/aa767916\(v=VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa767916(v=VS.85).aspx). Accessed on May 1, 2011.
- [2] Cisco Visual Networking Index: Forecast and Methodology, 2009-2014. [http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-481360\\_ns827\\_Networking\\_Solutions\\_White\\_Paper.html](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html). Accessed on May 1, 2011.
- [3] Dispelling the New SSL Myth. <http://devcentral.f5.com/weblogs/macvittie/archive/2011/01/31/dispelling-the-new-ssl-myth.aspx>. Accessed on May 1, 2011.
- [4] InternetProtocolInfo interface. [http://msdn.microsoft.com/en-us/library/aa767874\(v=VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa767874(v=VS.85).aspx). Accessed on May 1, 2011.
- [5] Network Simulator. [http://www.akmalabs.com/downloads\\_netsim.php](http://www.akmalabs.com/downloads_netsim.php). Accessed on May 1, 2011.
- [6] I. Akamai Technologies. Secure content deliver. [http://www.akamai.com/dl/feature\\_sheets/fs\\_edgesuite\\_securecontentdelivery.pdf](http://www.akamai.com/dl/feature_sheets/fs_edgesuite_securecontentdelivery.pdf).
- [7] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. RFC 4033: DNS Security Introduction and Requirements, 2005.
- [8] A. Barth, C. Jackson, and J. C. Mitchell. Securing Frame Communication in Browsers. In *Proceedings of the 17<sup>th</sup> USENIX Security Symposium*, San Jose, CA, July 2008.
- [9] R. J. Bayardo and J. Sorensen. Merkle Tree Authentication of HTTP Responses. In *Special Interest Tracks and Posters of the 14<sup>th</sup> International Conference on World Wide Web (WWW)*, Chiba, Japan, May 2005.
- [10] D. E. Bell. Looking Back at the Bell-LaPadula Model. In *Proceedings of the 21<sup>st</sup> Annual Computer Security Applications Conference (ACSAC)*, Tucson, AZ, Dec. 2005.
- [11] D. E. Bell and L. J. LaPadula. Secure Computer Systems: Mathematical Foundations. Technical Report ESD-TR-73-278, MITRE Corporation, Bedford, MA, Nov. 1973.
- [12] K. J. Biba. Integrity Considerations for Secure Computer Systems. Technical Report ESD-TR-76-372, MITRE Corporation, Bedford, MA, Apr. 1977.
- [13] U. Blumenthal and P. Goel. RFC 4785: Pre-Shared Key (PSK) Ciphersuites with NULL Encryption for Transport Layer Security (TLS), 2007.
- [14] T. Close. Petname Tool: Enabling Web site Recognition using the Existing SSL Infrastructure. In *W3C Workshop on Transparency and Usability of Web Authentication*, New York, NY, Mar. 2006.
- [15] K. DeGrande. CDNNetworks, September 2010. Personal communication.
- [16] P. Eckersley and J. Burns. Observatory for the SSLiverse, July 2010. <http://www.eff.org/files/DefconSSLiverse.pdf>.
- [17] FiddlerCore. <http://fiddler.wikidot.com/fiddlercore>.
- [18] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. RFC2616: Hypertext Transfer Protocol – HTTP/1.1, 1999.
- [19] S. Friedl. An Illustrated Guide to the Kaminsky DNS Vulnerability. <http://unixwiz.net/techtips/iguide-kaminsky-dns-vuln.html>.
- [20] C. Gaspard, S. Goldberg, W. Itani, E. Bertino, and C. Nita-Rotaru. SINE: Cache-Friendly Integrity for the Web. In *Workshop on Secure Network Protocols (NPSec)*, Princeton, NJ, Oct. 2009.
- [21] R. Gennaro and P. Rohatgi. How to Sign Digital Streams. In *Proceedings of the 17<sup>th</sup> Annual International Cryptology Conference on Advances in Cryptology (CRYPTO)*, Santa Barbara, CA, Aug. 1997.
- [22] S. Hanna, R. Shin, D. Akhawe, P. Saxena, A. Boehm, and D. Song. The Emperor’s New APIs: On the (In)Secure Usage of New Client-side Primitives. In *Web 2.0 Security and Privacy (W2SP 2010)*, 2010.
- [23] J. Hodges, C. Jackson, and A. Barth. Http strict transport security (HSTS), 2010. <http://tools.ietf.org/html/draft-hodges-strict-transport-sec>.
- [24] C. Jackson, A. Barth, A. Bortz, W. Shao, and D. Boneh. Protecting Browsers from DNS Rebinding Attacks. In *Proceedings of the 14<sup>th</sup> ACM Conference on Computer and Communications Security (CCS)*, Alexandria, VA, Oct. 2007.
- [25] V. Jirasek. Overcoming man in the middle attack on Strict Transport Security, August 2010. <http://blog.jirasek.eu/2010/08/overcoming-man-in-middle-attack-on.html>.
- [26] A. Langley, N. Modadugu, and W.-T. Chang. Overclocking SSL. In *Velocity: Web Performance and Operations Conference*, Santa Clara, CA, June 2010. <http://www.imperialviolet.org/2010/06/25/overclocking-ssl.html>. Accessed on May 1, 2011.
- [27] E. Lawrence. Fiddler Web Debugging Tool. <http://www.fiddler2.com/fiddler2/>. Accessed on May 1, 2011.
- [28] D. Mosberger and T. Jin. httpperf—A Tool for Measuring Web Server Performance. *Performance Evaluation Review*, 26(3):31–37, 1998.
- [29] S. Ramachandran. Let’s make the web faster. <http://code.google.com/speed/articles/web-metrics.html>.
- [30] C. Reis, S. D. Gribble, T. Kohno, and N. C. Weaver. Detecting In-Flight Page Changes with Web Tripwires. In *Proceedings of the 5<sup>th</sup> USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, San Francisco, CA, Apr. 2008.
- [31] E. Rescorla. RFC 2818: HTTP Over TLS, 2000.
- [32] E. Rescorla. and A. Schiffman. RFC2660: The Secure Hypertext Transfer Protocol, 1999.
- [33] J. Ruderman. Same Origin Policy for JavaScript. <http://www.mozilla.org/projects/security/components/same-origin.html>. Accessed on May 1, 2011.
- [34] S. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The emperor’s new security indicators. In *Proceedings of the 28<sup>th</sup> IEEE Symposium on Security and Privacy*, Oakland,

- CA, May 2007.
- [35] K. Singh, A. Moshchuk, H. J. Wang, and W. Lee. On the Incoherencies in Web Browser Access Control Policies. In *Proceedings of the 31<sup>st</sup> IEEE Symposium on Security and Privacy*, Oakland, CA, May 2010.
- [36] S. Stamm, B. Sterne, and G. Markham. Reining in the Web with Content Security Policy. In *Proceedings of the 19<sup>th</sup> International World Wide Web Conference (WWW)*, Raleigh, NC, Apr. 2010.
- [37] A. Stubblefield, A. D. Rubin, and D. S. Wallach. Managing the Performance Impact of Web Security. *Electronic Commerce Research*, 5:99–116, January 2005.
- [38] H. J. Wang, X. Fan, J. Howell, and C. Jackson. Protection and Communication Abstractions for Web Browsers in MashupOS. In *Proceedings of the 21<sup>st</sup> ACM Symposium on Operating Systems Principles (SOSP)*, Stevenson, WA, Oct. 2007.
- [39] H. J. Wang, C. Grier, A. Moshchuk, S. T. King, P. Choudhury, and H. Venter. The Multi-Principal OS Construction of the Gazelle Web Browser. In *Proceedings of the 18<sup>th</sup> USENIX Security Symposium*, Montreal, Canada, Aug. 2009.
- [40] J. Wang. A Survey of Web Caching Schemes for the Internet. *SIGCOMM Computer Communication Review*, 29:36–46, October 1999.
- [41] A. Wolman, G. M. Voelker, N. Sharma, N. Cardwell, A. Karlin, and H. M. Levy. On the Scale and Performance of Cooperative Web Proxy Caching. In *Proceedings of the 17<sup>th</sup> ACM Symposium on Operating Systems Principles (SOSP)*, Charleston, SC, Dec. 1999.