Resource Hints in HTML5:
A New Pandora's Box of Security Nightmares

N. Vlajic, X. Y. Shi, H. Roumani, P. Madani
Department of Electrical Engineering and Computer Science
York University, Toronto, Canada
vlajic@cse.yorku.ca, xueshi@my.yorku.ca, roumani@cse.yorku.ca, madani@cse.yorku.ca

ABSTRACT

[1] To date, much of the development in Web-related technologies has been driven by the users' quest for ever faster and more intuitive WWW. One of the most recent trends in this development is built around the idea that a user's WWW experience can further be improved by predicting and/or preloading Web resources most likely sought by this user, ahead of time. *Resource hints* is a set of features introduced in HTML5 and intended to support the idea of predictive preloading in the WWW. Unfortunately, as the very actualization and the present use of the *resource hints* have been almost exclusively driven by the speed and end-user experience in mind, the opportunities for their misuse in terms of other user-related metrics (user privacy and reputation, as well as business analytics) appear to be considerable.

In this article, we outline four different scenarios (i.e., attacks) in which the *resource hints* end up turning the browser into a dangerous tool that acts without the knowledge of and/or against its very own user. What makes these attacks particularly concerning is the fact that they are extremely easy to execute, and they do not require that any form of client-side malware be implanted on the user machine. While one of the attacks is (just) a new form of the well-known *cross-site request forgery* attacks, the other attacks have not been addressed much or at all in the literature. Through this work, we ultimate hope to make the wider Internet community critically rethink the way the *resource hints* are implemented and used in today's WWW.

CCS CONCEPTS

• **Security and privacy** → **System security**; *Browser security*
• **Security and privacy** → **Software and application security**; → Web application security

KEYWORDS

resource hints, unsolicited Web requests, user privacy, user reputation, browser forensics, Web attacks, HTML5, Chrome

1 INTRODUCTION

With the ever-growing importance and prevalence of WWW-based services and applications, we are becoming increasing reliant on the use and performance of Web browsers – software applications that allow users to access, traverse and retrieve the WWW resources. And, while in the past Web browsers were almost exclusively built for and used on desktop and laptop computers, nowadays any device capable of connecting to the Internet (e.g., mobile phones, smart watches, wearable tracking devices) are likely to host one or multiple Web browsers. In fact, the modern-day dilemma is not so much whether a Web browser should be available on an Internet-enabled device (regardless of it size and capability), but what can be done to make the performance of that browser faster and more user friendly.

Users' quest for ever faster and more intuitive WWW has been the driving force behind the evolution of Web-browser technology as well as numerous Web-related protocols. One of the most recent stages in this evolution is driven by the idea that a user's WWW experience can further be improved by predicting and/or preloading Web resources most likely sought by that particular user.

One specific mechanism that was recently introduced in order to make this idea of 'predictive preloading' possible is the so-called *resource hints* feature in HTML5. In particular, resource hints is a term that covers four different types of resource (pre)loading: preconnect, dns-prefetch, prefetch and prerender – all four being implemented as a relation (rel) type/attribute of HTML5's Link Element <link> [1]. When found in a Web-page (i.e., HTML5 document) resource hints are intended to instruct the browser to get hold of resources that are related to or are part of the most likely next-page

ARES '17, August 29-September 01, 2017, Reggio Calabria, Italy

ACM ISBN 978-1-4503-5257-4/17/08…$15.00
http://dx.doi.org/10.1145/3098954.3104046

navigation, ahead of time. Thus, if/when the user actually decides to request the given page, the respective resources will be simply pulled out from the 'background', giving an illusion of instantaneous (near zero-delay) retrieval.

Now, an average Web user is likely to consider the resource hints features useful, as they undoubtedly have the potential to facilitate faster browsing experience. As a result, in many browser types resource hints are enabled 'by default'. This - combined with the fact that users generally tend to keep the default settings of their applications unchanged [3] - further implies that the execution of resource hints is likely to be enabled in a significant number, if not the majority, of browser instances currently used in the Internet.

While we do not intend to question the practical usefulness of resource hints from the performance/speed point of view, the work presented in this paper seeks to address the potential negative implications of their use. Namely, the resource hints are generally designed to be executed without the user's direct involvement (i.e., knowledge or approval) and in an obscure 'behind the scenes' manner. And even though this un-intrusiveness has its obvious advantages when it comes to speed and convenience, it can also be easily misused - both intentionally and unintentionally - by turning a browser into a dangerous tool that acts without the knowledge of and/or against its very own user. The goal of our work is to bring awareness to these possibilities, and to make the wider Internet community rethink the way resource hints are implemented and used in today's WWW.

The reminder of this article is organized as follows. First, we discuss the significance and implications of using IP addresses as a means of identifying and tracking WWW users – a common Internet practice that is a precursor to the problem discussed in this article. Then, we provide a detailed overview of the four resource hints features/tags and present some of our experimental results concerning the execution of these tags in Google Chrome. Subsequently, we outline four different scenarios in which resource hints have the potential to negatively impact the user's security, reputation, and business operation. Finally, we close the article with conclusions and recommendations for future research.

2 BACKGROUND AND MOTIVATION

2.1 Relation Between a User and His Computer/Browser

As today's world grows ever more reliant on the WWW, the boundaries between humans and their respective Internet-enabled devices and browsers are becoming increasingly blurred. Namely, in many disciplines it has become a common practice to assume that a user's device and browser are nothing but a mere extension of the user, and their only mission is to carry out the tasks explicitly requested by the user. Consequently - in all but cases of a verifiable device/browser infection by a computer malware - the user may be considered fully accountable for actions or requests executed by their device/browser.

The concepts of user tracking and Web-related forensics are perhaps the best illustration of how tight the 'coupling' between users as persons and their device/browser is. For example:

- In user tracking, the IP address and cookies associated with a user's device (i.e., browser) are used to identify that particular user in the 'on-line world'. Subsequently, all observed Web requests that happen to carry those particular IP address and/or cookies are assumed to be generated with the full knowledge and intent of the given user and, as such, are used to track the user's online behavior as well as gauge their interest in different product and services [4]. User tracking mechanisms put relatively little (if any) effort in distinguishing between genuine user requests and those that were automatically generated by the user's browser.
- The goal of Web-related forensics is to gather information about which Web sites and files a user has accessed while browsing the WWW, in order to prove or disprove a claim of misconduct. The places where forensics-related artifacts are typically collected include: a) the browser history and cache on the user's device (if accessible), and/or b) the log files of the edge gateway that connects the user to the Internet, and/or c) the log files of the Web server(s) hosting the disputed files. If any evidence of the disputed files being accessed through the user's device/browser (while in the user's possession) is found in either a) or b) or c), the user himself could be held responsible – even without an explicit proof that the user, not the browser, was the one who actually initiated those requests.

The study presented in this article is motivated by the fact that the resource hints features outlined in the preceding section, when combined with our tendency to assume that devices and browsers are nothing but innocuous and trustworthy 'extensions' of their owners/users, can lead to a number of potential misuses. To lay a foundation for further discussion of this issue, we proceed by providing an outline of a typical WWW client-server architecture and its most significant elements and interactions as pertaining to our study.

2.2 Typical WWW Client-Server Architecture

The below figure outlines the most significant elements of a typical WWW client-server architecture, and those include:

a) The client, which in the case of the WWW is a Web browser running on the user' device. The device could be either 'fixed' (e.g., a desktop computer) or 'mobile' (e.g., a laptop, tablet or smartphone), and is uniquely identified either with a static IP address (common scenario in fixed enterprise networks) or a dynamic IP address (common scenario in cellular and public WiFi networks).

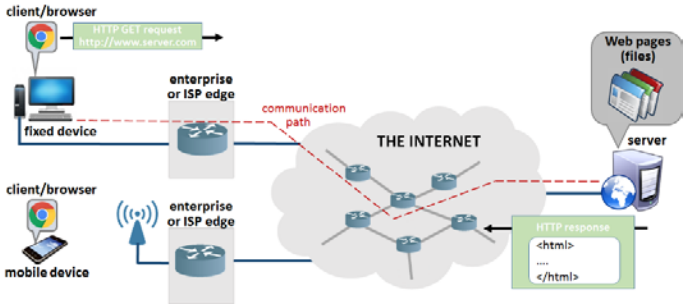

Figure 1: Typical WWW Client-Server Architecture

b) The edge network, which provides physical connectivity between the user device and the rest of the Internet. This could be either an enterprise edge network (e.g., when the device is used at work), or an ISP edge network (e.g., when the device is used at home). In either case, the edge network typically contains one or multitude of specialized devices which engage in monitoring and/or logging of the passing traffic (e.g., gateway routers, firewalls, proxies, …).

c) The Internet core, which is responsible for routing packets, including those that carry client-server HTTP requests and responses, from their source to the intended destination.

d) The server, which in the case of the WWW is a machine capable of hosting and sharing Web-pages (i.e., files) over the Internet, and typically performs continuous and detailed logging of all incoming traffic.

Now, whenever a particular client requests a Web-page from a particular server (by means of a GET HTTP request), various types of 'artifacts' related to this event get recorded at various points along the communication path between the two entities. For example:

i) On the client side, the URL of the requested page gets recorded in the browser history, while the resources that the requested page is made of get stored in the browser cache (once they actually arrive from the server). As earlier indicated, browser history and cache are of great significance from the perspective of Web forensics, since they can help prove that a particular Web request has taken place. Nevertheless, the main challenge of relying on browser history and cache as forensics evidence is that they are owned by and directly accessible to the user, and as such could be easily modified or deleted (intentionally or unintentionally) or simply rendered unavailable if the user decides to deny access (in which case a search warrant is required to be able to access these resources).

ii) The given HTTP request is likely recorded, together with the traffic of other users, in the logs of the specialized devices in the edge network (gateway, firewall or proxy). It should be noted, however, that edge networks are not always mandated to record these logs, hence from the forensics point they may have limited practical relevance.

iii) The intermediate routers in the Internet core could also keep a record of the given HTTP request in their own traffic logs. However, due to the high volumes of passing/recorded traffic, these logs are generally kept for a very short interval of time. Consequently, their practical use as forensics evidence is rather limited, similar to ii).

iv) The server logs is the final place where the given HTTP request gets recorded. In general, server logs have particularly important significance from the forensics point of view, for two main reasons. Firstly, most organization tend to retain their Web server logs over long periods of time. Secondly, in most organizations Web server logs are well protected and could only be altered by the site administrator. Hence, when a record of a Web request arriving from a particular client/host (i.e., IP address) is found in these logs, it is impossible to deny the authenticity of the given event - unless one can prove that the logs were altered (e.g.) by a malicious site administrator or some form of malware implanted on the server system.

With the above facts in mind, we further focus on the following fundamental question: for an HTTP request generated by the client/browser (e.g., while responding to a resource hint tag/command found in a rendered Web-page), is there a way of determining whether the given request was generated a result of an intentional action by the user, or perhaps it was generated without the user's knowledge and approval? Put another way, we are set to examine whether the artifacts collected along the given communication path provide enough information to tell these two different types of requests apart.

3 RESOURCE HINTS EXECUTION IN CHROME

In this section we provide a more detailed look at the four different types of resource hints mechanisms that can prompt a browser to perform various forms of resource preloading, without the user's explicit knowledge and intervention.

3.1 Resource Hints in HTML5

Hypertext Markup Language (HTML) is a well-known and widely used interpreted tagged markup language that enables creation of Web-pages (hypertext documents). In the most recent version of the protocol (HTML5) a special new set of features have been introduced in order to support the idea of 'instant' (zero-delay) Web-page load. Namely, as pointed in [6] and [7], a browser that starts downloading a Web-page only after the page has been explicitly requested by the user will inevitably result in substandard browsing experience that is riddled with various types of network delays. (These delays include: DNS lookup delay, TCP handshake delay, SSL negotiation delay, delay to obtain base HTML page … [5], [6].) The only way to spare the user from experiencing the browsing/network delays is by trying to anticipate their requests ahead of time, and then preload the most critical resources associated with those requests even before the actual 'click on the link' action occurs. That way, the resources

will be readily available when the user actually requests them, giving an illusion of an instantaneous (zero-delay) download.

Now, the idea of 'instant' browsing is not entirely new. This concept was originally supported through the implementation of Web-cache – a memory location where the resources of previously visited Web-pages are stored, allowing that these resources be instantaneously retrieved whenever the user decides to subsequently revisit them. Unfortunately, as such, Web-cache is of no use when it comes to retrieving new pages that have not been previously requested. To enable zero-delay browsing of pages that are to be visited for the first time, or pages that have been purged or expired from the cache, HTML5 has come up with a set of features commonly referred to as resource hints.

According to [1], there are four different type of resource hints provisions in HTML5.

a) dns-prefetch is a resource hint that can be used to suggest a browser to perform a DNS prefetch (i.e., IP lookup) for a particular hostname. The following is a situation where this feature might be useful in practice. Imagine the user is currently visiting page_A.html hosted on server_1.com, and there is a high likelihood that the Web-page the user is going to visit next is page_B.html located on another server (server_2.com) – as illustrated in Figure 2. To expedite the loading of page_B.html (if and when the user requests it), the below tag could be placed in the <head> section of page_A.html:

<link rel="dns-prefetch" href="//server_2.com">

That way, the browser would start performing the DNS lookup for sever_2.com right away (while the user is still viewing page_A.html), making sure that the IP address of server_2 is obtained even before the user actually clicks on http://server_2.com/page_B.html.

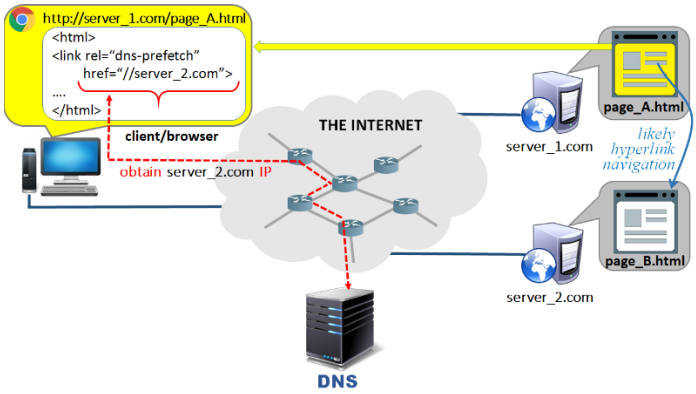

Figure 2: Linked pages hosted by different servers

b) preconnect is a resource hint option that can be used to initiate an early connection with a Web server, which includes the DNS lookup, TCP handshake, as well as optional TLS negotiation. As such, preconnect clearly goes step further in minimizing/masking networking delays relative to dns-prefetch.

In the example of Figure 2, the following tag placed in the <head> section of page_A.html would prompt the user's browser to establish an early (pre)connection with server_2.com.

<link rel="preconnect" href="//server_2.com">

Also, in the given example, the decision whether to use preconnect or just dns-prefetch for server_2.com should be closely tied to the actual probability that the user navigates to page_B.html from page_A.html. Clearly, the higher this probability, the more reasonable it would be to use the preconnect resource hints option.

c) prefetch is a resource hint option that further builds on the functionality of a) and b). Namely, in addition to performing the DNS resolution and establishing a connection with a particular server, prefetch also allows that some resources (e.g., the base HTML file of a Web-page, images, JavaScript-s, CSS-s, etc.) be downloaded from this server ahead of time and stored in the browser cache. For example, in the scenario of Figure 2, the following tag placed in the <head> section of page_A.html would prompt the user's browser to download and cache the base HTML file of page_B.html – the key Web resource (and the first one to be retrieved) during the rendering of this page.

<link rel="prefetch" href="//server_2.com/page_B.html">

Clearly, by allowing that whole parts of a page be obtained by the browser - even before the page gets actually requested - prefetch enables even further reduction in networking/browsing delays. However, given the communication and storage overhead associated with prefetch, it is recommended that this resource hints option be used only in cases when the probability that the user actually navigates to a specific page is greater than in the case of a) or b).

d) prerender is the most encompassing resource hints option - it allows not only that the base HTML file and all other components of a page get preloaded ahead of time, but also that the page itself gets fully laid out, its respective CSS-s applied and JavaScript-s executed. Put another way, it is as if the page is open in a hidden tab, and the moment the user navigates to the page's URL, the hidden tab is immediately swapped into view [5]. As such, prerender is the only resource hints option that can truly cut the browsing delay down to zero, giving an illusion of truly instantaneous browsing.

In the scenario of Figure 2, the following tag placed in the <head> section of page_A.html would prompt the user's browser to prerender (i.e., preload and preassemble) the entire page_B.html.

<link rel="prerender" href="//server_2.com/page_B.html">

Now, it should be pretty clear that out of all four resource hints options, the use of prerender is associated with the most significant communication, storage and processing overhead. Consequently, the use of this option should be reserved only for cases when the navigation to a specific page is highly probable if not absolutely certain.

The above suggestions are merely recommendations pertaining to the resource hints options in HTML5 as outlined by World Wide Web Consortium (W3C) [1]. Unfortunately, the actual implementation of the resource hints options in real-world browsers has neither been standardized nor mandated. As a result, there has been a significant variation in the number and actual implementation of different resource hints options by different browser types. (For more see [6], [7], [8]). Given that for the majority of Internet users Google Chrome happens to be the browser of choice [2], our discussion focuses on this particular browser type. Specifically, in the proceeding section, we present some of our experimental results pertaining to the behavior of Google Chrome when encountering different resource hints options in the browsed pages.

4 EXPERIMENTAL SET-UP AND RESULTS

In order to gain a better understanding of how Google Chrome deals with different HTML5 resource hints options when encountering them in a browsed page, we have built an experimental client-servers framework as outlined in Figure 3. The 'client' in this framework is the latest version of Google Chrome (Chrome v.57) running on a laptop PC. The 'server' is set up on the Amazon Cloud (http://ec2-54-186-72-100.us-west-2.compute.amazonaws.com) and hosts a repository of test Web-pages. We have chosen to code the pages of this repository in php instead of plain html in order to be able to prevent their caching on the client side, as well as to be able to implement and examine the general impact of cookies on pages referenced in resource hints tags.

The test pages of our framework are grouped into two sets. The pages of the first set are designed to be directly visible/accessible to the user, and each of them hides one resource hints option in its respective php/html code (A.php, B.php, C.php, D.php). The other set is comprised of pages referenced in the resource hints tags of the first set, and is not intended to be directly visible/accessible to the user (A_hidden.php, B_hidden.php, C_hidden.php, D_hidden.php). With this structure, if the pages of the second set - or their respective resources - ever get requested, that is a clear indication that the browser itself (not the user) has triggered those requests while processing the resource hints tags in the pages of the first set. (Note that, because of the way resource hints are intended to work as well as the way our framework is designed, requests for the pages of the second set not only get generated without the user's direct knowledge and involvement, but the user also never gets to know when those resources actually arrive at their browser.)

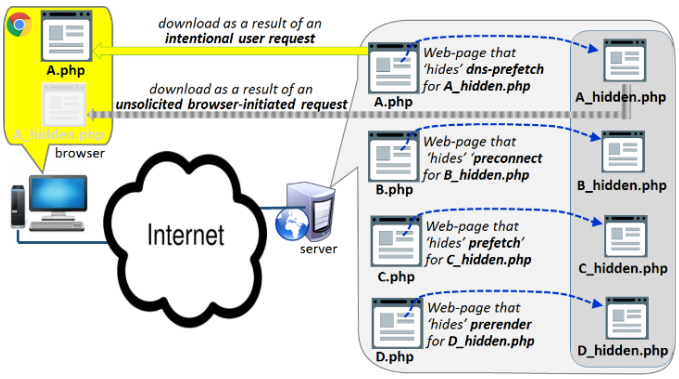

Figure 3: Experimental framework for evaluation of Chrome behavior when browsing pages with resource hints options

Table 1: Artifacts collected on client and server side when resource hint options found in a Web-page

	browser-side artifacts				server-side artifacts
resource hints option	**effect on Chrome history**	**effect on Chrome cache**	**effect on Chrome DNS cache**	**effect on cookies**	**server side log**
DNS-prefetch	*no effect*	*no effect*	*no effect*	*no effect*	*no GET request received at the server*
preconnect	*no effect*	*no effect*	*no effect*	*no effect*	*no GET request received at the server*
prefetch	*no effect*	*prefetched page/resource showed up in cache*	*showed up as a subresource of the calling web-site*	*cookies created*	*a GET request for prefetched resource/page received at the server (unless page/resource found in cache)*
prerender	*no effect*	*prerendered page showed up in cache*	*showed up as a standalone record (same as a user initiated visit)*	*cookies created*	*a GET request for prerendered page received at the server (unless page found in cache)*

In our experimentation, we first performed intentional requesting/retrieval of pages A.php to D.php (Figure 3) through the client - Chrome v.57 browser operating on a machine in our departmental network. Subsequently, we examined the collected artifacts pertaining to these requests both on the client and on the server side. The most significant of our observations are presented in Table 1, and can be summarized as follows:

1) The requesting of pages A.php and B.php (i.e., pages that contain DNS-prefetch and preconnect resource hint options in their respective HTML5 code) did not leave any permanent artifacts related to A_hidden.php and B_hidden.php - either on the client or on the server side. Such a result could have been expected, as these two particular resource hints options do not 'trigger' application-level preloading of resources referenced in their <link> tags. Instead, DNS-prefetch and preconnect facilitate only 'lower level' (DNS and TCP) domain-name resolution and connection set-up.

2) On the other hand, the requesting of pages C.php and D.php (i.e., pages that contain prefetch and prerender resource hint options in their respective HTML5 code), did leave a number of artifact related to C_hidden.php and D_hidden.php on the client and on the server side. In particular:

2.a) On the client side, both (prefetched) C_hidden.php and (prerendered) D_hidden.php were not only retrieved but also ended up being stored in the browser cache. Furthermore, a cookie associated with each of these pages was created and placed in the browser's cookie cache. Finally, a DNS record pertaining to both pages was stored in the browser's DNS cache. All in all, the way the browsers went about retreiving C_hidden.php and D_hidden.php was not much different from the way A.php to D.php were retrieved – even though the latter group of pages was explicitly requested by the user, while the user had no way of knowing that the former group of pages was ever requested and/or retrieved. (The only noticeable difference between the two groups is that the retrieval of A.php to D.php was recorded in the browser history, which was not the case for C_hidden.php and D_hidden.php.)

2.b) On the server side, HTTP GET requests for both C_hidden.php and D_hidden.php appeared in the server logs. More importantly, these two requests looked identical to the requests for pages A.php to D.php, in terms of their (HTTP) content. In other words, based on what was recorded in the sever logs, it was impossible to distinguish between the user's intentional requests - for A.php to D.php - and the requests that were issued automatically by the browser without the user's knowledge and approval (for C_hidden.php and D_hidden.php).

Following the experimentation with the framework outlined in Figure 3, we conducted another experimental study, where the Web objects referenced in A.php to D.php were pages hosted on another server. The observations concerning the recorded artifacts in this experiment were identical to the ones presented hereinabove (i.e., in Table 1).

Our experimentation also looked at the use of multiple prerender and prefetch tags inside the same Web-page. Our observation is that in case of multiple prerender tags in a Web-page, only one of these tags is executed at the time, while the respective (prerendered) page gets placed in the browser's RAM. (The likely reason why Chrome and other browser do not allow simultaneous prerendering of multiple pages is to prevent potential overloading of the browser's RAM, which would degrade the overall browser performance.) On the other hand, there seem to be no limit on the number of prefetch tags that get executed in a Web-page. Once retrieved, each of the prefetched resources ends up being stored in the browser's cache.

4 RESOURCE HINTS IMPLICATIONS ON USER PRIVACY, REPUTATION AND BUSINESS PERFORMANCE

In this section, we present four different scenarios in which resource hints are used as the main attack vector against a targeted Web user. The key characteristics of all four attacks is the fact they are extremely easy to execute, as they do not require that any form of client-side malware be implanted on the victim machine. The only precondition for their successful execution is to be able to lure the targeted user (victim) into visiting a specially crafted decoy Web-page. As indicated in [10], there are numerous well-known and very effective techniques which the attacker could deploy to lure a victim into visiting a decoy Web-page – ranging from various site-promotion techniques (e.g., in blogs and social media sites) to the use of targeted phishing emails.

Scenario 1: Framing attack.

The term 'framing attack' was introduced in [10], and it refers to a scenario in which false (digital) evidence is planted on the victim's computer, without requiring physical or remote access to their machine and without involving any form of client-side malware. The sole goal of this attack is to incriminate or discredit the victim in the context of their social, workplace, business or political life.

To provide an illustration of how a framing attack could be accomplished by means of HTML5 resource hints, imagine a situation where Trudy is a disgruntled employee working at a research company. Trudy holds a special grudge towards Bob – a manager that she directly reports to. As a form of revenge against Bob, Trudy decides to format one of her upcoming reports as an HTML5 document. Inside this document, she 'hides' several dozens (or more) of resource hint tags – each prefetching a highly inappropriate (e.g., child pornography or terrorism-related) Web-page. By means of JavaScript, Trudy also ensures that the execution of each prefetch tag occurs at a different point in time, thus mimicking the way a human user would go about retrieving a sequence of such Web-pages.

The 'reporting' day has come, and Bob opens the document that Trudy has referred him to. The (visible) content

of the document seems very relevant, and Bob spends quite some time viewing the document in his browser. Clearly, while Bob is reading the visible content, his browser (in the background) retrieves/prefetches the inappropriate pages (i.e., hidden resource hints) one-by-one, as illustrated in Figure 4. Bob, obviously, remains completely unaware that these downloads are taking place.

At the same or later point in time, the company's Web-content firewall generates an alert pointing to Bob's machine (i.e., his machine's IP) as the source of requests for inappropriate content. The company's authentication system verifies that the requests were generated while Bob was logged in and using the machine. From the forensics point of view, these pieces of evidence are often enough to 'points fingers' to Bob, and hold him accountable.

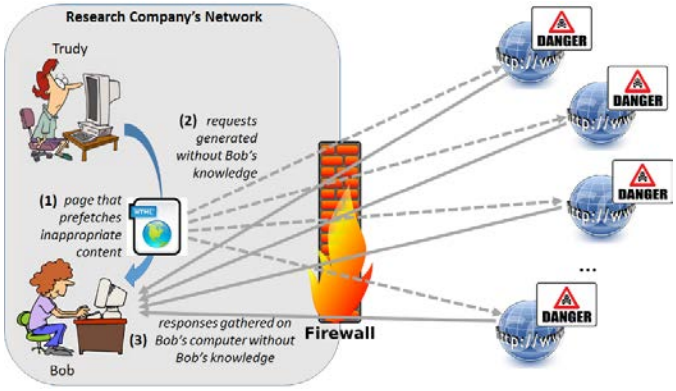

Figure 4: Framing attack

Now, depending on how severe the company's policy pertaining to inappropriate use of resources is, Bob could experience a whole range of possible outcomes – from receiving a simple warning to facing serious disciplinary actions and possibly termination. The only way Bob could avoid these repercussions and clear his name is by providing aggregate browsing-related artifacts from his computer (spanning over a period of time before and after the actual incident) to relevant authorities. While an adequate expert analysis of these artifacts could potentially succeed in putting 'all the pieces of the puzzle together' and identify the actual cause of the inappropriate requests, the implications on Bob's privacy could be significant - especially if Bob had used his own personal device (as in the case of BYOD) to view Trudy's page. In addition, by the very virtue of being linked with actions that are considered ethically and/or legally unacceptable, Bob is likely to experience unnecessary scrutiny with all the accompanying negative implications on his professional and personal life. (The best illustration of this are the cases of Julie Amero [12] and Michael Fiola [13]. These two people, in two different instances, were wrongly charged with downloading of child pornography. In both cases it was ultimately proven that the downloading of the inappropriate material was caused by malicious software and their respective names were cleared. Still, as stated by both people, the conducted trials have had lasting negative effect on their lives as well as the lives of their families.)

According to our knowledge, [10] is the only other research work that, in addition to introducing, has also studied the actual mechanisms of executing a framing attack. The idea specifically suggested in [10] is similar to the one outlined in Figure 4, except that the obscure/decoy requests are not generated via resource hint tags (prefetch of prerender) but instead by means of two better known and more widely used HTML tags - <iframe> and <img>. However, as indicated in [10], for these framing attacks to actually be successful, the attacker needs to take extra measures towards 'obscuring' the objects/Web-pages referenced in the decoy <iframe> and <img> tags (i.e., make sure that they go unnoticed by the victim once they are retrieved/rendered by the browser). Possible approaches to ensuring that the decoy <iframe> and <img> objects remain 'invisible' include:1) minimizing their size to 0x0 pixels, 2) hide them under another overlaid iframe/image, 3) make them invisible through CSS (e.g., by setting their display attribute to none). It should be noted, though, that the same object obfuscation techniques are required and deployed by many other types of browser-based attacks, such as clickjacking and cross-site request forgery. These specific attacks have been around for more than a decade, and as a result, the majority of today's Web-vulnerability scanning tools (e.g., Burp [14]) are programmed to spot and block Web-pages suspected of object obfuscation. Consequently, a framing attack based on the use of <iframe> and <img> decoy tags (as proposed in [10]) could potentially be detected and prevented by these tools. On the other hand, a framing attack based on the use of HTML5 resource hints (as suggested in this work) would virtually go unnoticed by these same scanning tools. Namely, while a 'malicious' <iframe> and <img> could be detected (i.e,. labeled as such) by looking for signs of obfuscation, there are no clear mechanisms or indicators which could help in distinguishing between a benign and a malicious <prerender> or <prefetch> tag. (Recall, the very purpose of these tags is to facilitate 'invisible' preloading of Web object. Furthermore, the objects are supposed to remain hidden until explicitly requested by the user.)

Scenario 2: Targeted DoS Attacks.

Now, imagine that in the previously depicted story, instead of tarnishing Bob's reputation within their organization, Trudy decides to execute her revenge by affecting the 'outside' reputation of Bob's machine (i.e., IP address), with the ultimate goal of having Bob's IP address blacklisted and denied service.

In particular, imagine that Trudy knows of a Web-site that Bob likes to frequently visit, such as the Web-site of his bank or a specific news-agency Web-site. In that case, Trudy could hide a very large number of prefetch references targeting this particular Web-site (its various pages/resources) inside her 'malicious' Web-page, as illustrated in Figure 5. As many other

similar organizations, Bob's bank is likely to perform comprehensive intrusion detection monitoring of the incoming Web traffic, in order to spot and blacklist all misbehaving users. Given that the avalanche of requests coming from Bob's machine is very reminiscent of a denial of service (DoS) attack, it is quite possible that Bob's IP would end up on the bank's blacklist, at least for a period of time. Consequently, during that period of time, even Bob's legitimate requests would be rejected (since coming from the same IP), and Bob would be cut off from the online services of his bank. (We refer to this attack as 'targeted DoS', as it ensures that one specific user is denied service by one particular Web-site.)

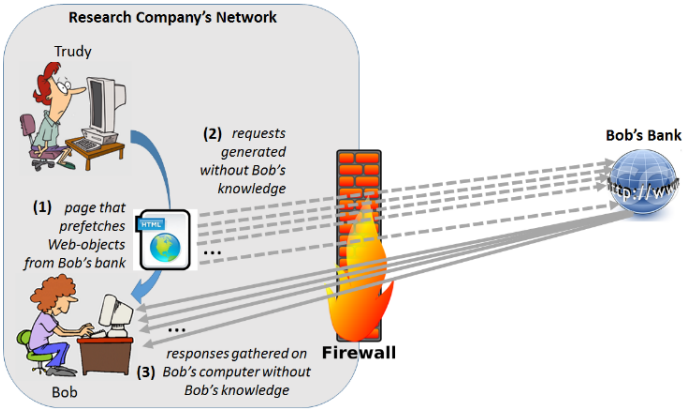

Figure 5: Targeted DoS attack

Scenario 3: Cross-Site Request Forgery Attack.

Cross-site request forgery (CSRF) is a well-known type of attack that occurs when a malicious Web-site causes a user's browser to perform an unwanted action on a trusted site for which the user is currently authenticated [15]. More specifically, CSRF attack requires that the user first gets successfully authenticated to a legitimate Web-site (e.g., by means of cookies), as illustrated in Figure 6. If following that the user visits a malicious Web-page (shown in Figure 7), the malicious page can force the user's browser to make unsolicited request towards the site for which the user is currently authenticated. By default, the browser will attach the legitimate previously set cookie(s) to each of the unsolicited/malicious requests, which will make the server's job of distinguishing between genuine user requests and those that were triggered by the malicious Web-page hard if not impossible.

CSRF attack have been traditionally accomplished by 'hiding' the unsolicited HTTP requests of the malicious page inside <img> and <iframe> HTML tags, as in the case of the framing attack described in [10]. However, we have already explained that many of today's Web-vulnerability scanning tools are capable of detecting such 'basic' variants of CSRF attacks, simply by looking for signs of <img> or <iframe> obfuscation. Consequently, from the attacker's point of view, hiding the unsolicited CSRF HTTP requests inside <prefetch> and <prerender> tags is a viable and far more lucrative alternative, as today's Web-vulnerability scanning tools generally do not look for signs of misuse in any of the four resource hints options.

Figure 6: User authentication by means of cookies

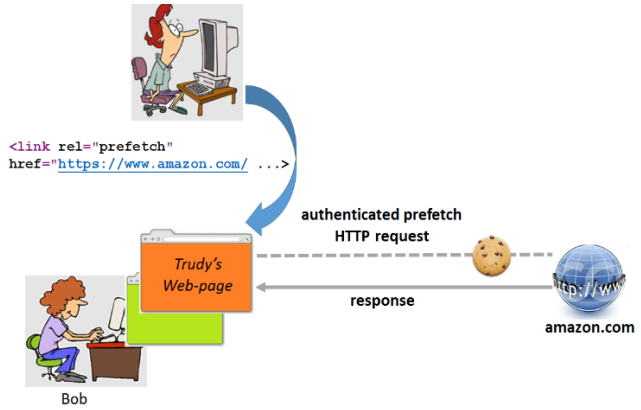

Figure 7: CSRF attack following user authentication

(According to our knowledge, no previous work has looked at the use of resource hints options in the context of CSRF attacks. Our group has recently conducted a study on the feasibility of CSRF attacks on amazon.ca and ebay.ca by means of resource hints, which resulted in the discovery of open vulnerabilities in both sites. The findings of this study are currently under submission to [16].)

Scenario 4: Data-Analytics Pollution Attack.

(This particular type of attack has also not been previously discussed in the literature. Its main aim is to impact the performance of an on-line business by distorting its Web-site based data analytics.)

As the premise for this attack, we imagine Trudy to be the owner of a small business, and Alice to be her direct business competitor. Both businesses have online presence which is critical for the success of their operation. Namely, not only that the two businesses advertise and sell their products through their respective Web-sites, but they also heavily rely on the Web (server-log) analytics to better understand where their customers come from and what they are looking for.

In order to 'pollute' the logs of Alice's Web server, and thus negatively impact Alice's business intelligence, Trudy has come up with the following plan: In the Web-page(s) of her own Web-site, Trudy has hidden numerous prerender and prefetch tags

referencing various (strategically chosen) pages from Alice's Web-site. Thus, whenever Trudy's customers visit her Web-site, their respective browsers end up generating a slew of 'polluting' requests towards Alice's Web server – see Figure 8. Obviously, because of the way the resource hints are intended to work, Trudy's customers will be completely unaware that their browser has participated in a 'data polluting' attack. At the same time, the performance of Trudy's Web-site will remain entirely unimpacted by the attack, as the retrievals of prerender/prefetch resources from Alice's Web-site will always take a lower priority and occur only during the browser's idle times.

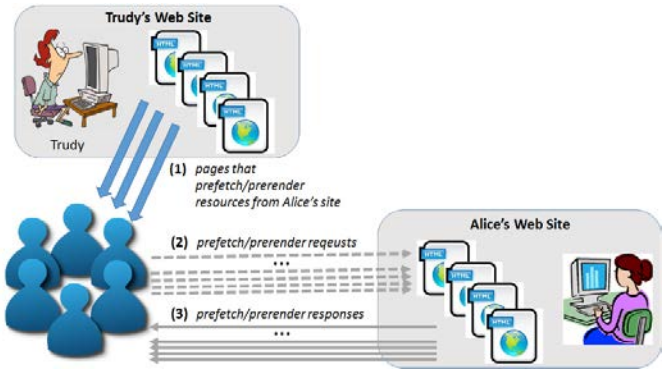

Figure 8:Data polluting attack

As for Alice's ability to detect this attack and identify all the polluting requests – the only piece of information that she possibly could rely on is the referer field in the incoming HTTP requests. (Referer field identifies the address of the Web-page from which the user/browser has accessed, or moved to, the current Web-page.) In the case of Trudy's attack, this field would be referring to the pages of her Web-site, thus indirectly reviling the true origin of the 'polluting' requests. Though, if Trudy wanted to make her attack particularly stealthy, she could implement the following meta tag in the <head> section of her Web-pages used to refer to resources in Alice's Web-site:

<meta name="referrer" content="none">

That way, HTTP requests arriving to Alice's Web-site by means of Trudy's Web-pages would not contain any referral data. Consequently, Trudy's attack would remain virtually undetectable.

6 CONCLUSIONS AND FUTURE WORK

The goal of this article is to bring awareness to a slew of vulnerabilities that have been created with the introduction of HTML5 resource hints. We have provided examples of specific threats and attacks that are easy to mount and can have serious implications.

In order to mitigate these risks, further work is warranted and it can be structured within the general framework of handling threats; namely, to deter and block, and failing that, to be able to recover from an attack. These can be achieved by a combination of one or more of the following measures:

1. Browsers should have an option to disable resource hints so users can block potential attacks. Chrome provides such an option but is set to "allow" by default.
2. Browsers should make resource hints transparent, so that users are aware of them, without impacting the user experience.
3. Discriminating browser-initiated loads from user-initiated ones is currently done through the HTTP Purpose header, which is not logged by most servers. We propose that this be elevated to a request parameter (i.e., ?purpose=prefetch) so that it can be readily available during forensics investigations.
4. Increase awareness, particularly amongst expert witnesses and analysts, of the footprint left by resource hints. For example, if a page appears in a browser's cache but not in its history is a telltale sign that this was not a deliberate user-initiated retrieval.

We plan to pursue some of these directions in future works.

REFERENCES

[1] Resource Hints. W3C Working Draft, 27 May 2016. DOI= https://www.w3.org/TR/resource-hints/
[2] StatCounter Global Stats. Top 5 Desktop, Tablet & Console Browsers from Aug 2015 to Aug 2016. DOI= http://gs.statcounter.com/?PHPSESSID=oc1i9oue7por39rmhqq2eouoh0
[3] C. Arthur. Why the default settings on your device should be right first time. theguardian.com, December 2013. DOI= https://www.theguardian.com/technology/2013/dec/01/default-settings-change-phones-computers
[4] M. Bichler. The Future of eMarkets: Multi-Dimensional Market Mechanisms. Cambridge University Press, 2001.
[5] February 2016 Web Server Survey. Netcraft, February 2016. DOI= https://news.netcraft.com/archives/2016/02/22/february-2016-web-server-survey.html
[6] Ilya Grigorik. High Performance Networking in Google Chrome. January 2013. DOI= https://www.igvita.com/posa/high-performance-networking-in-google-chrome/
[7] Ilya Grigorik. High Performance Browser Networking. O'Reilly, 2013.
[8] B. Jackson. Resource Hints – What is Preload, Prefetch and Preconnect?. KeyCDN Blog. July 2016. DOI= https://www.keycdn.com/blog/resource-hints/
[9] W3Tech Web Technology Surveys. Usage of Cookies for Websites. September 2016. DOI= https://w3techs.com/technologies/details/ce-cookies/all/all
[10] N. Gelernter, Y. Grinstein, A. Herzberg. Cross-Site Framing Attacks. 31st Annual Computer Security Applications Conference (ACSAC'15). Los Angeles, CA, USA. December 2015.
[11] G. Rydstedt, E. Bursztein, D. Boneh, C. Jackson. Busting Frame Busting: a Study of Clickjacking Vulnerabilities on Popular Sites. IEEE Symposium on Security and Privacy (S&P'10). Oakland, California, May 2010.
[12] PCWorld. The Julie Amero Case: A Dangerous Farce. Dec 2008. DOI= http://www.pcworld.com/article/154768/julie_amero.html
[13] The Register. How malware frames the innocent for child abuse. Nov 2009. DOI= https://www.theregister.co.uk/2009/11/09/malware_child_abuse_images_frame_up/
[14] Burp. DOI=https://portswigger.net/vulnerability-scanner/
[15] Cross-Site Request Forgery (CSRF) Prevention Cheat Sheet. OWASP, March 2017. DOI= https://www.owasp.org/index.php/Cross-Site_Request_Forgery_(CSRF)_Prevention_Cheat_Sheet
[16] A. Basit, N. Vlajic. CSRF Attack Using HTML5 Resource Hints: A New Face of an Old Enemy. Submitted to 2017 IEEE Cyber Science and Technology Congress.