

# Security Vulnerabilities in Modern Web Browser Architecture

Marin Šilić, Jakov Krolo, and Goran Delač

Faculty of Electrical Engineering and Computing, University of Zagreb

Unska 3, 10000 Zagreb, Croatia

Phone: +385 1 6129 549

E-mail: marin.silic@fer.hr, jakov.krolo@fer.hr, goran.delac@fer.hr

**Abstract** - The Web today has become the most used and popular platform for application development. In the beginnings of the Web, applications provided users just the ability to browse and read content. The expansion and adoption of new web technologies has led to a significant increase in development and, more importantly, usage of web applications that allow users to create their own content and impact their life (e.g. e-banking, e-commerce, social networks). Web 2.0 applications introduced new possibilities for both users and application developers, but also created new security concerns. Almost every Internet user uses a web browser to access any content on the Internet. Each web application is designed and developed to be executed inside the web browser. The web browser mediates between users and applications. In such an architecture, malicious applications could be loaded and executed inside the web browser, making it a vulnerable point in preserving security. Modern web applications demand for a new web browser architecture design that will meet new security requirements arisen with Web 2.0. In this paper, we study web browser vulnerabilities, analyze popular web browsers architecture and present how they cope with potential security threats.

## I. INTRODUCTION

In its basis, the Web was designed for browsing static web pages and reading content. With the recent technological improvements, the Web has become a platform for application development. The turning point was the invention and adoption of AJAX technology [1] which turned the old concept of static web pages into the novel approach of deploying interactive web applications. AJAX and similar web application development technologies, often collectively referred to as Web 2.0 technologies, led to the creation of a variety of numerous worldwide-oriented web applications. Contemporary web applications like e-banking, e-commerce, social-networking sites, blogs, and video-sharing sites provide users not just the ability to view information and access content, but also the ability to contribute and create their own Web content, express their creativity and share knowledge and information.

The nature of Web 2.0 applications requires users to provide their identity and private data like user-names, passwords, credit card numbers, mailing addresses, social security numbers, etc. Those applications are designed to be executed inside the web browser, which is a mediator between users and applications. Web browser exploits have a bigger impact than ever before, and thus web browser designers have to pay more attention to security than ever before. Knowing the security holes in a web browser, attackers can create malicious web applications in order to compromise other users' security.

Many different web applications can be executed simultaneously within the web browser. Some applications can provide users with useful functionality, while other can be malicious with the sole purpose of compromising security. Regardless of their purpose, each application in the browser has its own security settings which define access privileges and rights for the user's local file system. For example, a web application should be allowed to access local file system in order to upload a certain file only with the user's explicit approval. On the other hand, browsers have a local storage where user's sensitive data like passwords, cookies, bookmarks, browsing history, temporary files, and cache are stored. Modern browsers need to assure that a web application can not access that storage, and can only retrieve private data (e.g. cookies) related to that same application.

In order to protect the user, some browsers enforce a strict security policy, which isolates applications inside the browser by their origin and does not allow access to subresources from other origins. Such a restrictive policy would require architectural restructuring of the existing Web. On the other hand, users expect browsers to be compatible with the existing Web architecture and render their popular applications properly. The desirable goal in browser design is to achieve user protection and still provide compatibility with existing web applications.

The majority of modern browsers still use the original monolithic architecture design [2]. Monolithic browser architecture has many disadvantages that concern client code execution. Failure caused by a single web application crashes the entire browser instead of just the application that caused it. In terms of better user experience, the user should be able to interact with other running web applications. From the aspect of security, if the browser as a vulnerable monolithic structure gets compromised, an attacker could execute arbitrary code with user's privileges and rights. Such a security lapse could cause severe damage on the local machine. In spite of this issue, modern applications require a browser architecture that provides both browser security and compatibility with the existing Web architecture. That can be achieved with modular browser architecture where, in contrast to monolithic one, each application is executed in its own sandbox with restricted privileges.

The rest of the paper is structured as follows: Section II explains modular browser architecture and compares it with the monolithic one. In section III, Google Chrome browser is reviewed, as an implementation example of the modular browser architecture. We analyze how Chrome responds to major browser security threats. Section IV describes related browsers based on modular architecture and compares them with Chrome. Finally, the concluding remarks are given in Section V.

## II. MULTI-PROCESS BROWSER

As a result of recent a breakthrough in the Web technology utilization, contemporary web applications behave more like complex programs rather than simple documents for browsing. Most of current web browser architectures are still for the most part monolithic, usually designed for tasks like browsing and rendering of static web pages. Monolithic architectures do not provide enough isolation between concurrently executed web programs, so the execution often ends in misbehavior caused by the lack of security, fault-tolerance, memory management or degraded performance. Early operating systems had the same program isolation issues. MS-DOS and MacOS allocated single address space which sometimes caused programs to interfere with each other. Modern operating systems isolate each program in its own separate process. Similarly, modern browsers should isolate web programs and modularize their execution. Modularization can be achieved by assigning each web program to a specific operating system process within the browser.

### A. Monolithic Browser Architecture

Figure 1 shows monolithic web browser architecture, most common for current web browsers. In monolithic architecture, all web browser components are placed in a single operating system process. The HTML Renderer component parses each page code and generates a DOM tree. Document Object Model (DOM) tree is a web page representation that can be accessed and modified by the script code. To run the script code, browser utilizes the JavaScript Engine component.

Monolithic architecture has a lot of disadvantages that impact user experience, fault-tolerance, accountability, security, memory management and performance. Regarding user experience and fault-tolerance, any web program or browser component that encounters a crash, takes down the complete web browser. To cope with this issue, some browsers, with a monolithic architecture, provide a reload feature to aid the recovery process after the crash. But still, as a result of the crash user might lose valuable data (e.g. unsaved email drafts, e-banking transactions, purchase orders) stored as JavaScript state in the memory. Considering accountability, monolithic architecture provides resource usage statistics for the entire web browser. However, the web program responsible for a poor performance of the entire browser can not be identified in a monolithic architecture. Another disadvantage of monolithic browser architecture is memory management. Browser process in OS is a long life process compared to the life of web programs that are executed in the browser. A web program running in the browser might require a lot of memory allocation and could contain memory leaks. This can result in a large and fragmented memory space that is allocated to the browser process. Once the web program terminates, the memory still remains large and fragmented which can ultimately lead to poor performance. As far as performance is concerned, monolithic architectures can cause resource demanding web programs

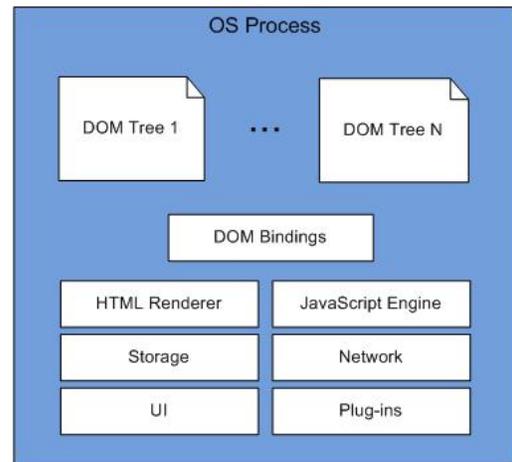


Figure 1. Monolithic browser architecture

to compete for CPU resources. Also, monolithic architecture can block a browser UI thread because web program is running actions, like executing a synchronous *XMLHttpRequest*. Both of this causes user-perceived delays on UI level and lower performance for the entire browser. Security of monolithic architectures rely entirely on the browser components logic to completely isolate different web programs and prevent any information flow between web objects in different web programs. However, security holes in browser design or implementation create opportunities for malicious web programs and attackers to install malware, steal files or access private data and compromise user's security.

Despite all its disadvantages, monolithic architecture has been preserved in the majority of web browsers. This is mainly because it is difficult and challenging to isolate web programs in the browser and still enforce browser compatibility with the Web architecture. One naive approach in accomplishing this goal is to isolate each web page in the browser in its own process. Unfortunately, this would degrade browser's ability to run many popular applications like sites that use pop-up windows or embed content in a separate frame from a different location. Another approach could be to isolate web programs by their origins. However, sometimes pages with different origins need to communicate with each other and sometimes pages with the same origins are not related at all.

### B. Modular Browser Architecture

Figure 2 presents web programs isolation model, implemented in Google Chrome web browser, based on open source Chromium project [3]. The key point in modular browser design is to isolate web programs, but provide compatibility with the current Web. Presented model introduces ideal abstractions: web program and web program instance.

*Web program* is a set of connected web pages including all their subresources that are required to provide certain functionality. For example, the iGoogle page contains the parent page, script libraries, images, and gadgets sites embedded in their frames. Since browsers allow users to

visit multiple instances of the same page (e.g. user can open two iGoogle pages in different browser tabs), the web program instance abstraction is introduced. A *web program instance* is defined as a set of pages from a web program that are connected in the browser and allowed to access and manipulate each others content. The web program abstraction is implemented as a *site*, while web program instance abstraction is implemented as a *site instance*.

Browsers allow related pages to communicate by enforcing *Same Origin Policy (SOP)* [4]. SOP conducts access control based on the page origin, which includes protocol, full host name, and port of each page. Pages with the same origins are grouped together and allowed to manipulate each others content. Page subresources can be included from some other origins, but their origin is considered same as the origin of the enclosing page. If origins do not match, pages are mainly isolated. Origin does not provide enough distinction among pages to define a site because a web page can change its origin dynamically. The origin can be changed within a limited range, from sub-domain to more general domain and only up to the *registry controlled domain name* [5], which is the most general part of the host name before suffix (e.g. *zemis.fer.hr* can be changed to *fer.hr*). Site is defined as a set of web pages with origins within the specific origin range, limited with protocol and registry controlled domain name.

Sometimes pages with a different origin are connected and share a communication channel. This happens in cases when a page opens content in a new window. The opener page then keeps a reference to a new window and the opened page can access the opener using the property *window.opener*. For example, the Gmail chat window opens in a new window when the conversation is initiated. The second case is when a page embeds content from a different origin in a separate frame. For example, the iGoogle page contains multiple gadget pages loaded in separate frames. The top window can access its frames using the *windows.frames* property, and each gadget can access the parent window using the *window.parent* property. Connections between pages that share a communication channel are kept as long as the parent browser window is alive. Even if the user navigates to another page or opens a new page in a new tab or window, references among those pages are kept. Chrome isolation model [3] defines the term *browsing instance* as a set of connected windows or frames that keep a reference to each other.

Site instance, the implementation of web program instance abstraction, is defined as a set of connected pages belonging to the same site within the browsing instance. All pages from the same browsing instance can reference each other's windows, but only the pages from the same site instance can access each other's DOM contexts. Figure 2 shows an example of two browsing instances. First browsing instance *B1* contains two site instances: *Sa* and *Sb*. Second browsing instance *B2* contains site instance *Sa*. Site instance *Sa* contains page *Pa*, while site instance *Sb* contains page *Pb*. Each page may contain subresources loaded from foreign origin (e.g. page *Pa* contains subresources *B* and *C*). Site instance *Sa* from *B1* and *Sa* from *B2* belong to the same site, but do not reference each other, although SOP would allow them to communicate.

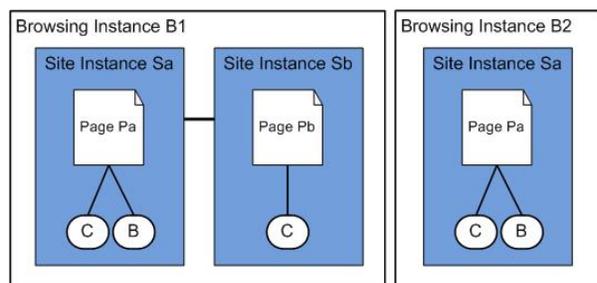


Figure 2. Isolation model in Chrome

On the other hand, site instances *Sa* and *Sb* from *B1* belong to different sites, they reference each other, but are still not allowed to manipulate each other DOM according to the SOP.

One approach to accomplish isolation is to assign each site instance to a single operating system process. In this case, there would be too many processes allocated for web browser execution. Another approach, used by Chrome, is to assign each browsing instance to a single operating system process. This way Chrome manages to isolate web programs and modularize web browser execution. Each web program is running in its own operating system process.

In comparison to monolithic architecture, modular architecture is superior considering user experience, fault-tolerance, accountability, security, memory management and performance. Considering user experience and fault-tolerance, each web program that crashes does not effect the execution of other running programs. In modular architecture each program performance can be easily monitored, thus modular architecture is superior in accountability. Memory management in modular architecture is effectively conducted; each program has its own process and allocated memory. Once the program finishes, memory is released and can be assigned to some other program. The fact that each program has its own process assures better performance. Modular architecture leaves scheduling issues to the OS, so web programs can run in parallel. A deeper insight into security aspects of modular architecture is given in the following section.

### III. SECURITY OF CHROME

The definition of web program isolation model is used as a base for modular architecture implementation. In this section we present modular browser architecture implemented in the Chrome browser. Furthermore, we categorize the main security treats for modular browser architecture and analyze security aspects of the Chrome browser in relation to those threats.

#### A. Chrome Architecture

Architecture of the Chrome browser [6] is given in Figure 3. Chrome consists of three different modules: rendering engine, browser kernel and plug-ins. Each of these modules is isolated in its own operating system process. Rendering engine converts HTTP responses into rendered bitmaps, browser kernel interacts with the OS,

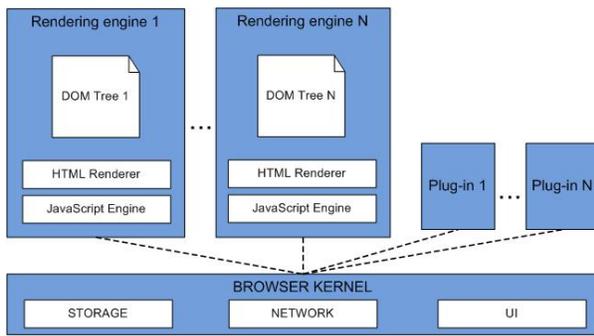


Figure 3. Chrome architecture

while the plug-ins module is responsible for execution of each plug-in.

The *rendering engine* runs in a sandbox with restricted privileges, deprived of access to the OS resources. Each isolated web program in the browser is assigned to its own rendering engine. The rendering engine is responsible for parsing web content, creating DOM tree representation in the memory, manipulating the DOM tree while executing script instructions. Also, the rendering engine enforces SOP and interacts directly with untrusted web content. Historically, most of the web browser security vulnerabilities were detected in the parsing and decoding tasks. Thus, the rendering engine, as an isolated process, does most of the parsing like HTML, CSS, XML, JavaScript, regular expressions parsing and image decoding. To interact with the user and OS, the rendering engine utilizes simple and restricted browser kernel APIs.

The *Browser kernel* runs with full user privileges on behalf of the user. It manages each instance of the rendering engine and implements browser kernel APIs. The browser kernel is responsible for storage management because such an activity requires file system access. By accessing the storage, browser kernel comes in contact with sensitive data, like cookies, bookmarks, passwords, etc. Furthermore, the browser kernel is responsible for executing network operations, e.g. downloading images sending them to a rendering engine for decoding. Also, browser kernel interacts with the OS, handles user inputs and forwards them to the rendering engine assigned to the focused window. In addition, browser kernel keeps the information about privileges granted to each rendering engine, such as a list of files that certain rendering engine may upload.

*Plug-ins* run in their own process, independent from the rendering engine and browser kernel. Web compatibility requires plug-ins to run outside the sandbox, because plug-ins may require access to a microphone, web cam or local file system. Thus, plug-ins can not be placed inside the rendering engine since rendering engine runs in a sandbox. Plug-ins could be placed within the browser kernel, but in this case, a crash in plug-ins would take down the entire browser. Furthermore, a security hole in plug-in design or implementation could be exploited to compromise security and run arbitrary code with full user's privileges.

## B. Security Aspects

Similar to the Threat Model proposed by Mozilla [7], we categorize the main threats to modular browser architecture and discuss how Chrome responds to them. We discuss system compromise, data theft, cross domain compromise, session hijacking, and user interface compromise.

The *system compromise* threat refers to malicious arbitrary code execution with full privileges on behalf of the user. The majority of browser's vulnerabilities, concerning this threat, are detected in the rendering engine, while minority is related to the browser kernel and plug-ins. The compromised rendering engine runs within a Windows sandbox, with a Windows restricted security token, unlike the browser kernel which runs with Windows user's security token [8]. Although Windows sandbox restricts rendering engine to communicate with the OS, there are potential issues that can be exploited in order to compromise the system. For example, sandbox does not perform security token check if the sandboxed process is accessing the FAT32 file system. Most of the existing devices use NTFS file system, but some USB devices still use FAT32 formatting. In this scenario, compromised rendering engine could read and write the content on the USB drive. Also, Windows sandbox does not perform security token check or requires OS handle when a low-level privilege process attempts to open TCP/IP socket. However, these issues are flaws in the Windows sandbox rather than the Chrome architecture. Besides compromising the rendering engine, threats can also occur if the browser kernel gets compromised. Browser kernel compromising usually happens when it is being tricked from a compromised rendering engine. Browser kernel could perform unauthorized network or system task as a lack of parameters validation while executing operations invoked by compromised rendering engine through its APIs. Finally, the system could get compromised through the third component as well – plug-ins module. An attacker could do damage to the system if it exploits vulnerabilities in vendors' plug-ins, because they run outside the sandbox by default.

*Data theft* refers to the ability to steal local network or system data. This often happens in case a compromised rendering engine requires uploading or downloading a file. In Chrome architecture, the rendering engine runs in a sandbox and has no direct access to the local file system. When uploading a file, the rendering engine uses browser kernel API for file upload. The browser kernel shows the upload file picker window and remembers which file was selected. This action is considered as an explicit user authorization to the associated rendering engine to upload that particular file. Such authorization lasts for the lifetime of the associated rendering engine. In the following step, the browser kernel uploads the file to the site instance running in the associated rendering engine. Also, when downloading files, the rendering engine utilizes browser kernel API. Since the download is initiated by the user, the browser kernel is authorized to download the resource from the specified URL. Some malicious site may include subresources with URLs that use the file scheme (*file:///* instead of *http://*). Chrome architecture prevents the rendering engine to issue network tasks, like requesting a resource from a specified URL. The rendering engine uses the browser kernel API to include subresources, and the

browser kernel then analyzes the resource URL and downloads the resource. Most of the rendering engines are not allowed to include subresources from URLs that use the file scheme. However, local files can be viewed in the Chrome browser through a dedicated rendering engine.

*Cross domain compromise* happens when a code originating from one *fully qualified domain name (FQDN)* [9] executes code in the context of, or reads data from, another FQDN domain without permission. One such attack is XML eXternal Entity (XXE) attack [10], in which the attacker's XML document, hosted at `http://attacker.com/`, includes an external entity from a foreign origin. For example, the malicious XML document might contain an entity from the `https://bank.com` or `file:///etc/passwd`. If vulnerable to XXE attacks, the browser will retrieve the content from the foreign origin and incorporate it into the attacker's document, making him able to read the content. Chrome, like many other browsers, uses *libXML* to parse XML documents. However, the architecture of Chrome is designed in such a way that it delegates parsing tasks to a sand-boxed rendering engine. The rendering engine does not prevent content being retrieved from foreign origins, but it rather passes the requests to the browser kernel. If the external entity URL was a web URL, the browser kernel services the requests. However, if the external entity URL was from the user's file system (i.e. from the file scheme), then the browser blocks the request, preventing the attacker from reading confidential information such as passwords. Chrome's modular architecture with sand-boxed rendering engines does not completely defend against the XXE vulnerability because the attacker is able to retrieve URLs from foreign web sites. To block such requests, the browser kernel would need to sacrifice compatibility with the Web architecture (e.g. ban cross-site images).

The threats that involve *session hijacking* compromise the session token by stealing or predicting a valid session token to gain unauthorized access to the honest web server.

Cross-site scripting (XSS) and cross-domain request forgery (CSRF) have become the two most scaled attacks regarding session hijacking. According to The Open Web Application Security Project (OWASP), those two kinds of attacks have been marked as No2 and No5 top security risks for web applications for the year 2010 [11]. In short, XSS exploits the client's trust of the content received from the server (by just sending text-based attack scripts that exploit the interpreter in the browser). This allows attackers to execute scripts in the victim's browser which can hijack user sessions, deface web sites, or redirect the user to malicious sites. A CSRF attack tricks (via image or script tags) a logged-on victim's browser to send a forged HTTP request, including the victim's session cookie and any other authentication information, to a vulnerable web application. This allows the attacker to perform any action on a vulnerable web server the victim is authorized to use. Chrome's architecture does not protect an honest web site if the site contains XSS or CSRF vulnerabilities. Chrome expects these sites to repair their vulnerabilities. The only helpful security measure Chrome has is the support for *HttpOnly* cookies, which can be used as a partial mitigation for XSS [12].

*User interface compromise* threat refers to the ability to trick the user into making incorrect trust decision, or directly provide confidential data using script UI manipulation. A popular attack that manipulates the UI is

*clickjacking* [13]. In clickjacking, an attacker loads a subresource from some other origin and places it as transparent content in front of the visible content. The user thinks he is clicking on the objects he sees, but actually he clicks on the transparent content. Furthermore, user interface compromise also refers to implementing annoying scripting actions like hogging the CPU or memory, moving dialogs faster than the user can respond, and opening dialogs in an endless loop. From today's perspective, no modern browser provides enough security restrictions to defend itself from scripting disruptions. Introduction of any limitations that are related to window manipulation or disabling pop-ups, provides less functionality and options for site developers and lowers compatibility with the Web. However, Chrome uses limitations on windows manipulation [14] with scripting languages such as: taking full screen, specifying screen dimensions and position, hiding the URL and status bars.

#### IV. RELATED ARCHITECTURES OVERVIEW

Some of the known web browsers that are based on the modular architecture are OP [15], Tahoma [16], and Gazelle [17]. They are not widely popular due to their main disadvantage of not being fully compatible with the Web architecture. This means that many web applications are not usable with those browsers due to their lack of support for some web techniques modern web applications are based on. However, these browsers introduce many interesting concepts regarding security and could possibly lead to a better security model for the future browser architectures. Such architectures would still have to be compatible with the Web in order to be widely acceptable.

*OP* web browser introduces a modular architecture that consists of the following browser components: UI, web page, storage, network and browser kernel. Each component runs in its own operating system process. Browser kernel runs with full privileges and behaves like the operating system micro-kernel. It coordinates the communication among browser components using message passing mechanism. Compared to Chrome, proposed architecture provides strong isolation among web sites and a higher level of protection, but also disables some popular web features like inter-frame communication, file uploads and downloads. These characteristics make the OP browser incompatible with many popular web sites. In addition, the OP browser enforces restrictive plug-ins security policy. The architecture does not allow plug-ins to run with full privileges in their own process on the whole browser level. Instead, plug-ins run within the web page component with restrictive privileges. Plug-ins are allowed to access the resources that correspond to the origin of the site whence the plug-in object is embedded. Sandboxing plug-ins using restrictive policy provides higher protection, but makes the browser less compatible with the current Web.

*Tahoma* architecture introduces a new concept for web application execution. Each web site is running on its own virtual machine within the protected framework named browser operating system (BOS). BOS manages each virtual machine's network and UI tasks. Each virtual machine manages its own storage, cookies, bookmarks, history and has no access to the user's local file system. There is a strict isolation among different running virtual

machines. Tahoma architecture introduces new possibilities for web application execution. Since a web application is running on virtual machine, web application developer can deploy application written in machine code. Each web site owner should create manifest file for his web site. That manifest file contains information about the site, like the list of URLs site is communicating with or weather the site uses machine code or standard HTML renderer. When first visiting the site, a user receives site manifest and needs to approve the site before the execution begins. Tahoma architecture is revolutionary and provides a high level of protection. However, this architecture is completely incompatible with the current Web and requires current Web architecture restructuring.

The *Gazelle* browser architecture contains a browser kernel and rendering engine process, similar to Chrome. Chrome places resources within the same renderer according to the registry controlled domain name policy, while *Gazelle* places resources within the same renderer according to the SOP. *Gazelle* architecture provides stronger isolation that concerns inter-frame scripting. In case the web page embeds content in a separate frame, *Gazelle* places parent and child frame into a different renderer processes and allows them to communicate using limited browser kernel API. On the other hand, Chrome places the parent and child frame in the same renderer process, but the communication among them is restricted according to the SOP. Cross scripts and style sheets are placed within the same renderer process both in Chrome and *Gazelle*. However, because of the cross-scripting and inter-frame communication limitations, *Gazelle* is not quite compatible with the current Web. For, example, *Gazelle* does not allow the frame to change its *document.domain* property, which was essential for inter-frame communication prior to *postMessage* event introduction. Also, *Gazelle* introduces opaque display policy, which disallows cross-site content to be transparent and overlap the host site. This policy enhances the overall browser security and reduces UI manipulation, but still it is not quite compatible with the current Web architecture. *Gazelle* browser tends to protect different web sites from each other, while Chrome focuses on host machine's and user's protection.

## V. CONCLUSION

In this paper we described the concept of web program isolation in the browser as a response to new security challenges and performance demands introduced with the Web evolution in recent years. In addition, we compared the new modular architecture to the monolithic architecture, mostly used in current web browsers, and showed that modular architecture is superior in many important aspects.

For the purposes of this paper, the Google Chrome modular browser architecture has been chosen as a referent architecture model. We analyzed Chrome's behavior with respect to the most popular web browser security threats. The analysis showed that modular architecture of Chrome mitigates most serious treats that are related to system compromise and data theft. However, Chrome's architecture does not provide full protection. Threats related to cross-site attacking, session hijacking and user interface compromise are not quite mitigated. We reviewed

similar architectures implemented in OP, Tahoma and *Gazelle* web browser. These architectures sacrifice compatibility with the current Web architecture in order to provide higher level of security than Chrome.

## REFERENCES

- [1] L.D. Paulson, "Building Rich Web Applications with Ajax," *Computer*, vol. 38, no. 10, pp. 14-17, 2005
- [2] M. Andreessen, and E. Bina, "NCSA Mosaic: A Global Hypermedia System", *Internet Research*, vol. 4, no. 1, pp. 7-17, 1994
- [3] C. Reis and S.D. Gribble, "Isolating web programs in modern browser architectures", *Proceedings of the 4th ACM European conference on Computer systems*, April 01-03, 2009, Nuremberg, Germany
- [4] C. Jesse Ruderman, "The Same Origin Policy", [https://developer.mozilla.org/En/Same\\_origin\\_policy\\_for\\_JavaScript](https://developer.mozilla.org/En/Same_origin_policy_for_JavaScript)
- [5] Mozilla, "Public Suffix List", <http://publicsuffix.org/>, 2007
- [6] A. Barth, C. Jackson, C. Reis, and Google Chrome Team, "The Security Architecture of the Chromium Browser", *Technical report*, Stanford University, 2008 <http://crypto.stanford.edu/websec/chromium/chromium-security-architecture.pdf>
- [7] Mozilla Wiki: "High Level Threat Model for Process Isolation", <https://wiki.mozilla.org/Security/ProcessIsolation/ThreatModel>
- [8] Microsoft: "Restricted Tokens", February 2010, [http://msdn.microsoft.com/en-us/library/aa379316\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa379316(VS.85).aspx)
- [9] Indiana University Knowledge Base: "Fully qualified domain name", October 2009, <http://kb.iu.edu/data/aiuv.html>
- [10] G. Steuck: "XXE (Xml eXternal Entity) attack", October 2002, <http://www.securiteam.com/securitynews/6D0100A5PU.html>
- [11] The Open Web Application Security Project: *OWASP Top10-2010 rc1*, "The Ten Most Critical Web Application Security Risks", [http://www.owasp.org/images/0/0f/OWASP\\_T10\\_-\\_2010\\_rc1.pdf](http://www.owasp.org/images/0/0f/OWASP_T10_-_2010_rc1.pdf)
- [12] Microsoft: "Mitigating cross-site scripting with HTTP-only cookies", <http://msdn.microsoft.com/en-us/library/ms533046.aspx>
- [13] M. Balduzzi, M. Egele, E. Kirda, D. Balzarotti, C. Kruegel: "A Solution for the Automated Detection of Clickjacking Attacks", *ASIACCS'10*, Beijing, China, 2010
- [14] M. Zalewski: "Browser Security Handbook", 2009, <http://code.google.com/p/browsersec/wiki/Main>
- [15] C. Grier, S. Tang, S. T. King: "Secure web browsing with the OP web browser", *2008 IEEE Symposium on Security and Privacy*
- [16] R. S. Cox, J. G. Hansen, S. D. Gribble, H. M. Levy: "A Safety-Oriented Platform for Web Applications", *2006 IEEE Symposium on Security and Privacy*
- [17] H. J. Wang, C. Griery, A. Moshchukz, S. T. Kingy, P. Choudhury, H. Venter: "The Multi-Principal OS Construction of the *Gazelle* Web Browser", *MSR Technical Report MSR-TR-2009-16*