Researching Polymorphic Images for XSS on Google Scholar

Researching Polymorphic Images for XSS on Google Scholar - Lorenzo Stella, blog.doyensec.com.

- Published: date not stated
- Original: <https://blog.doyensec.com/2020/04/30/polymorphic-images-for-xss.html>
- Preserved from: https://blog.doyensec.com/2020/04/30/polymorphic-images-for-xss.html (stored) on 2026-08-11
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

Researching Polymorphic Images for XSS on Google Scholar · Doyensec's Blog

Researching Polymorphic Images for XSS on Google Scholar

30 Apr 2020 - Posted by Lorenzo Stella

A few months ago I came across a curious design pattern on Google Scholar. Multiple screens of the web application were fetched and rendered using a combination of `location.hash` parameters and XHR to retrieve the supposed templating snippets from a relative URI, rendering them on the page unescaped.

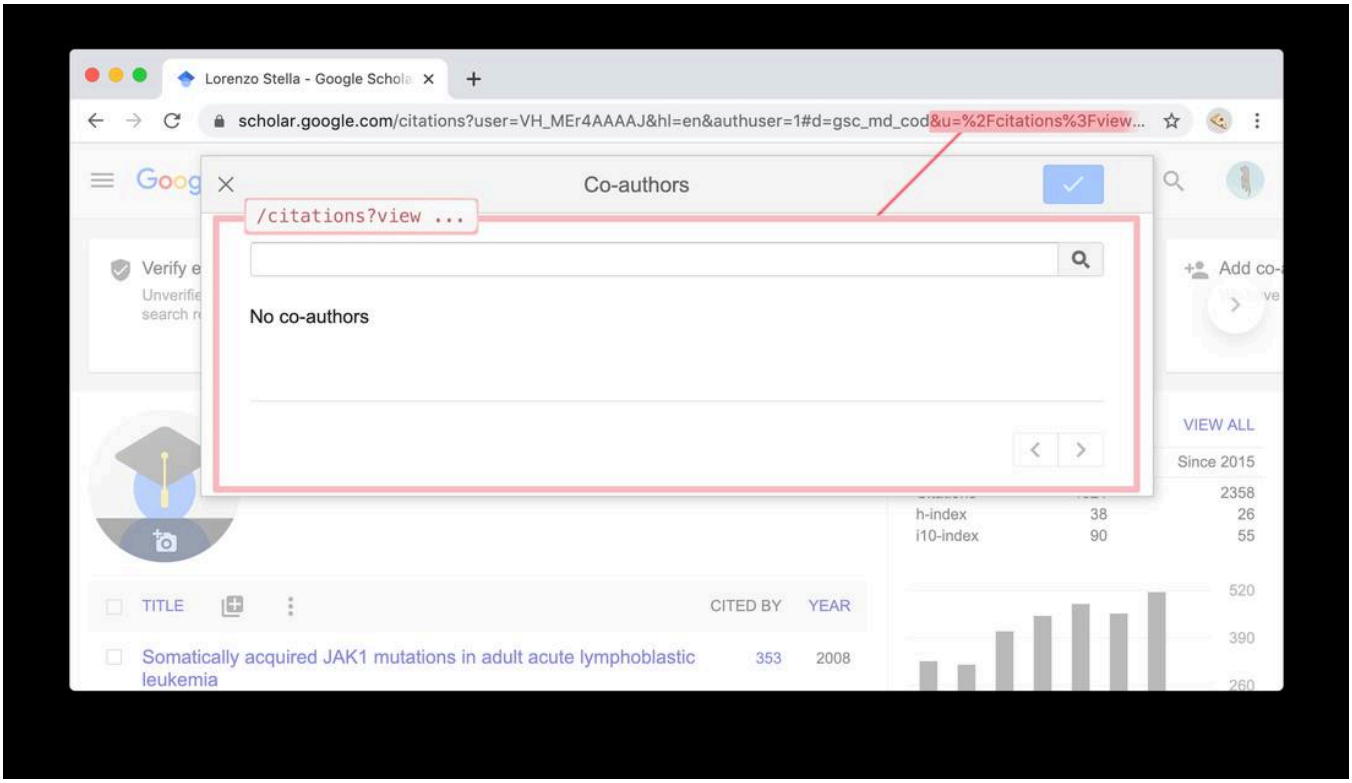

This is not dangerous per se, unless the platform lets users upload arbitrary content and serve it from the same origin, which unfortunately Google Scholar does, given its image upload functionality.

While any penetration tester worth her salt would deem the exploitation of the issue trivial, Scholar's image processing backend was applying different transformations to the uploaded images (i.e. stripping metadata and reprocessing the picture). When reporting the vulnerability, Google's VRP team did not consider the upload of a polymorphic image carrying a valid XSS payload possible, and instead requested a PoC||GTFO.

Given the age of this technique, I first went through all past "well-known" techniques to generate polymorphic pictures, and then developed a test suite to investigate the behavior of some of the most popular libraries for image processing (i.e. Imagemagick, GraphicsMagick, Libvips). This effort led to the discovery of some interesting caveats. Some of these methods can also be used to conceal web shells or Javascript content to bypass "self" CSP directives.

Payload in EXIF

The easiest approach is to embed our payload in the metadata of the image. In the case of JPEG/JFIF, these pieces of metadata are stored in application-specific markers (called `APPX`), but they are not taken into account by the majority of image libraries. Exiftool is a popular tool to edit those entries, but you may find that in some cases the characters will get entity-escaped, so I resorted to inserting them manually. In the hope of Google's Scholar preserving some whitelisted EXIFs, I created an image having 1.2k common EXIF tags, including CIPA standard and non-standard tags.

While that didn't work in my case, some of the EXIF entries are to this day kept in many popular web platforms. In most of the image libraries tested, PNG metadata is always kept when converting from PNG to PNG, while they are always lost from PNG to JPG.

Payload concatenated at the end of the image (after 0xFFD9 for JPGs or IEND for PNGs)

This technique will only work if no transformations are performed on the uploaded image, since only the image content is processed.

As the name suggests, the trick involves appending the JavaScript payload at the end of the image format.

Payload in PNG's iDAT

In PNGs, the iDAT chunk stores the pixel information. Depending on the transformations applied, you may be able to directly insert your raw payload in the iDAT chunks or you may try to bypass the resize and re-sampling operations. Google's Scholar only generated JPG pictures so I could not leverage this technique.

Payload in JPG's ECS

In the JFIF standard, the entropy-coded data segment (ECS) contains the output of the raw Huffman-compressed bitstream which represents the Minimum Coded Unit (MCU) that comprises the image data. In theory, it is possible to position our payload in this segment, but there are no guarantees that our payload will survive the transformation applied by the image library on the server. Creating a JPG image resistant to the transformations caused by the library was a process of trial and error.

As a starting point I crafted a "base" image with the same quality factors as the images resulting from the conversion. For this I ended up using this image having 0-length-string EXIFs. Even though having the payload positioned at a variable offset from the beginning of the section did not work, I

found that when processed by Google Scholar the first bytes of the image's ECS section were kept if separated by a pattern of `0x00` and `0x14` bytes.

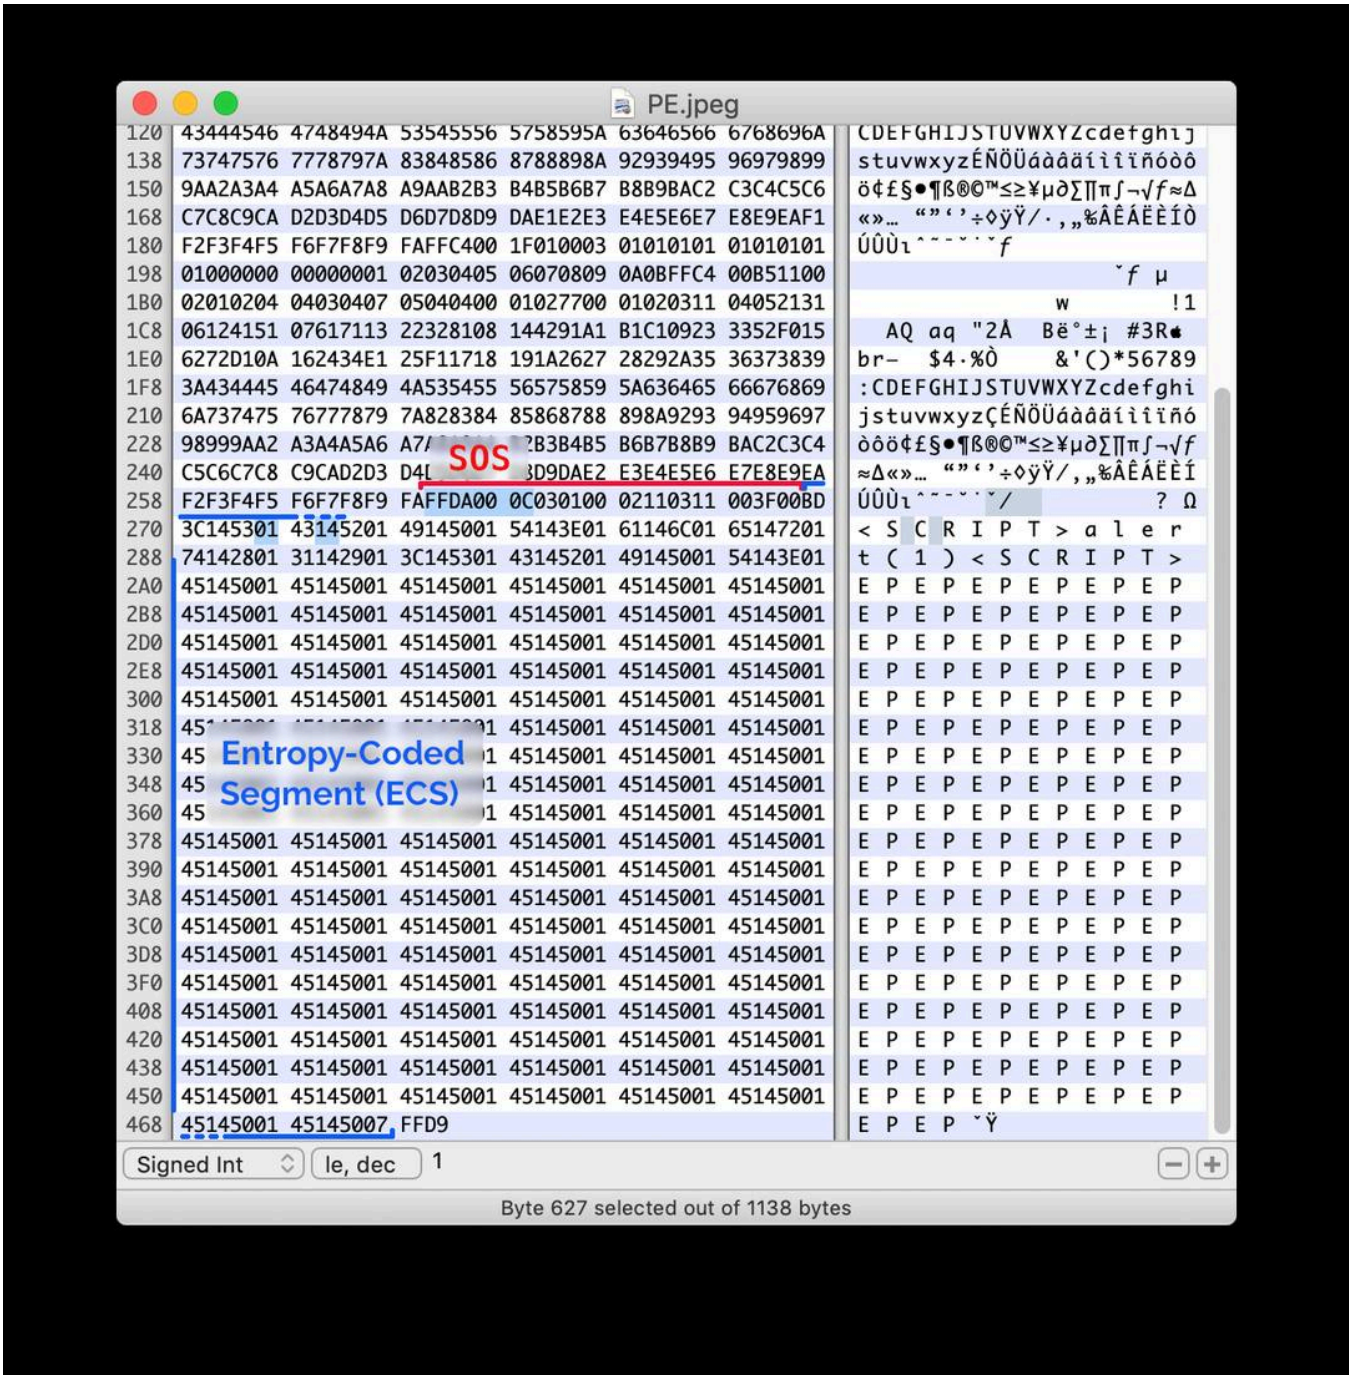

From here it took me a little time to find the right sequence of bytes allowing the payload to survive the transformation, since the majority of user agents were not tolerating low-value bytes in the script tag definition of the page. For anyone interested, we have made available the images embedding the onclick and mouseover events. Our image library test suite is available on Github as doyensec/StandardizedImageProcessingTest.

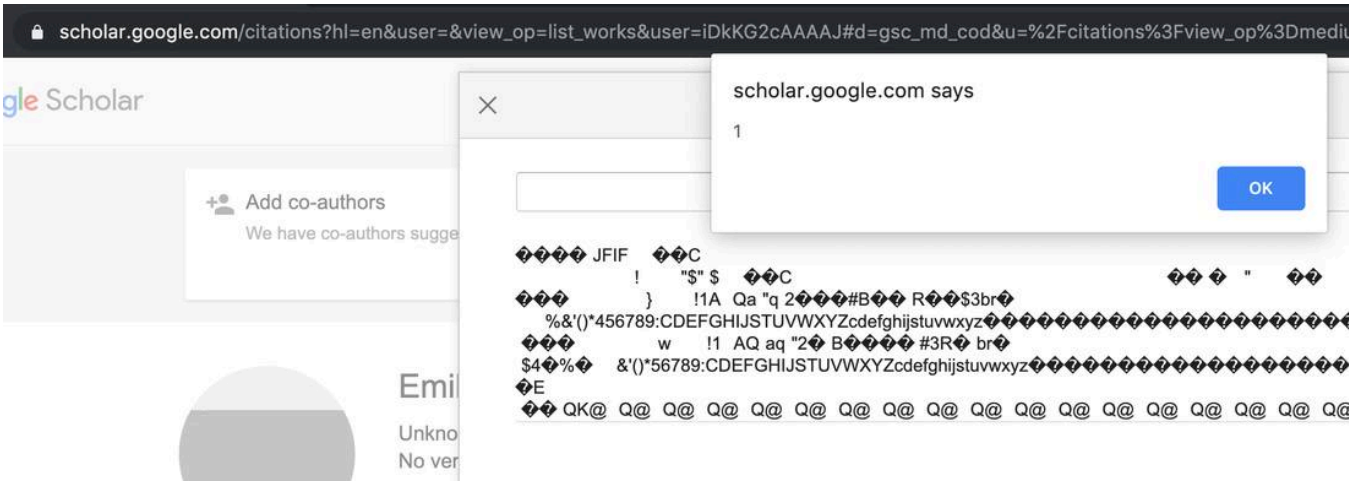

Timeline

- **[2019-09-28]** *Reported to Google VRP*
- **[2019-09-30]** *Google's VRP requested a PoC*
- **[2019-10-04]** *Provided PoC #1*
- **[2019-10-10]** *Google's VRP requested a different payload for PoC*
- **[2019-10-11]** *Provided PoC #2*
- **[2019-11-05]** *Google's VRP confirmed the issue in 2 endpoints, rewarded $6267.40*
- **[2019-11-19]** *Google's VRP found another XSS using the same technique, rewarded an additional $3133.70*
