

Protecting a Web Application Against Attacks Through HTML Shared Files

Francisco Corella

November 2008

Patent Pending

Abstract

Many Web applications have a file-sharing feature that allows Web users to share files by uploading them to, and downloading them from, a Web-accessible file repository. Shared files may include HTML files and other files containing scripts that are executed by the browser in the security context of the user that downloads the file. This opens the door to a range of cross-user attacks, including attacks by former users and even attacks by a user of a virtual application instance against a different virtual instance of the same application. Such attacks are in essence XSS attacks, but the usual defenses against XSS are typically not available, because shared files cannot be sanitized.

This paper proposes a countermeasure that Web applications can use against attacks through HTML shared files, without sanitizing those files. The countermeasure leverages the same-origin policy by the use of carefully tailored hostnames for serving user files and application pages, including two different hostnames for downloading a shared file, linked by a redirection step. Authentication is achieved by the use of different cookies for shared files and application pages, and ephemeral file-retrieval sessions.

1. Introduction

Many Web applications have a file-sharing feature that allows Web users to share files by uploading them to, and downloading them from, a Web-accessible file repository. (I will refer to files stored in the repository as *shared files* or *user files*.) Some of these applications provide file-sharing as their only purpose, while others provide file-sharing as one of several collaboration tools or social-networking features.

These applications typically have multiple *virtual application instances*, each having its own file repository and its own group of instance users who are allowed to access the file repository. (Applications with multiple virtual instances are sometimes referred to as *multitenant applications* [\[1\]](#), where the word *tenant* refers to the group of users of a virtual instance. In this paper I will simply say *application instance* to refer to a virtual application instance.)

Sharing files has the obvious risk that a virus or other malware may be present in a file being obtained through file sharing. This risk however is not specific to the fact of sharing; it exists whenever a file is obtained by any means. This risk is mitigated to a certain extent by warnings issued by browsers when downloading executable files, and may be further mitigated through the use of virus scanners. This paper is concerned, however, with a somewhat different and more insidious risk, one that is not mitigated by browser warnings or scanners.

Some of the files in an instance repository may be HTML files. When a user downloads an HTML file from the repository, the file may be displayed by the browser. This is a useful feature, as it allows one or more HTML shared files to serve as a private Web site accessible only to instance users, perhaps containing links to non-HTML files stored in the repository. But an HTML file displayed by the

browser may contain client-side code, such as scripts written in JavaScript or ActionScript, Java applets, or ActiveX controls. A script planted by a malicious user in a shared file, when downloaded by a victim user, runs in the security context of the victim user, and can take advantage of any authentication cookies set in the victim's browser to impersonate the victim and take advantage of any access privileges granted to the victim by the application.

(Similar considerations apply to other types of files that may contain scripts that are executed by the browser when a file is downloaded, such as SWF files or PDF files. References to HTML files in the rest of the paper should be understood to apply more broadly to any type of file that may contain client-side code, if the browser executes such code without warning and the effects of the code are subject to the same-origin policy as discussed below.)

This problem is a variation on the cross-site scripting (XSS) [2] problem faced by sites, such as social networking sites, that allow users to upload HTML content that is displayed to other users. Those sites typically defend against such attacks by sanitizing uploaded HTML content to remove or disable any scripts. Web applications with a file sharing feature, on the other hand, do not take precautions against attacks through shared files. Why is this?

One can think of two reasons. First, it is not a good idea to sanitize uploaded files. Users do not expect their files to be tampered with, and they may have legitimate reasons to share files that contain scripts. Second, an argument can be made that the Web application is entitled to assume that users of an application instance trust each other, and thus that it is not the responsibility of the application to prevent attacks of a user against another user (*cross-user attacks*).

The second reason, however, is not a valid one. The application has a responsibility to defend users against attacks through HTML shared files because those attacks can be used to defeat security mechanisms implemented by the application and negate security assumptions that users take for granted. For example, they can defeat user authentication and allow a user to impersonate another user, or a former user of an application instance to retain access to the instance without authorization. They can defeat the system of file permissions, if one is in place, giving a user access to files he or she is not authorized to download. They can even allow a user of one instance to access a different instance for which he or she does not have credentials (*cross-instance attack*).

This paper proposes a countermeasure that Web applications can use to protect users against attacks through HTML shared files, without sanitizing those files. The countermeasure leverages the same-origin policy [3] by the use of carefully tailored hostnames for serving user files and pages constructed by the application (*application pages*); two different hostnames linked by a redirection step are used to download a user file. Authentication is achieved by the use of different cookies for shared files and application pages, and ephemeral file-retrieval sessions.

The rest of the paper is organized as follows. [Section 2](#) surveys previous work. [Section 3](#) discusses attacks through HTML shared files in more detail. [Section 4](#) describes the proposed countermeasure. [Section 5](#) discusses possible variations on the countermeasure and secondary benefits. References can be found in [Section 6](#).

2. Prior Work

Little work has been done concerning the risk of cross-user attacks through HTML shared files. Application service providers (ASPs) do not discuss the risk, and some do not authenticate access to user files. US-CERT discusses the risks of file sharing in a Cyber Security Tip [4] but only considers peer-to-peer file sharing. Carnegie Mellon's MySecureCyberspace Web site discusses File Sharing security [5], but only considers peer-to-peer file sharing and centralized file sharing via FTP. A recent

thread in the Web Security mailing list [6] touched on Web-based file sharing security, but was mostly concerned with the general problem of malware being present in uploaded files.

An interesting proposal [7] aimed at mitigating XSS is applicable to HTML file sharing. It would allow the Web application to forbid the execution of scripts in files being downloaded, using a *Content-Restrictions* HTTP header. However the proposal requires buy-in from the browsers, and thus is not immediately usable. Also, users may want scripts to be executed, which the countermeasure of Section 4 below allows.

The only published thought that comes close to the proposal of Section 4 seems to be a single sentence in a Web page [8] that discusses XSS:

If you must allow unsanitized, untrusted HTML to be part of your site, ensure that those pages are not on the same hostname as where other users log in. (webmail, web hosts, attachments in a bug-tracking system, Google cache) (see Gerv's proposal)

(In the quoted sentence, "Gerv's proposal" is presumably a reference to [7].) This recommendation does not explicitly apply to file sharing, but it does contain the idea of protecting against malicious scripts by using different hostnames for different purposes. On the other hand it does not explain how to construct those hostnames nor how to provide authentication for accesses to both user files and application pages.

3. Cross-user and cross-instance attacks through HTML shared files

As briefly mentioned in the introduction, scripts embedded in shared files can be used to carry out a variety of cross-user and cross-instance attacks that cannot be mitigated by browser warnings or virus scanners.

The basic attack mechanism is as follows. The attacker, a malicious user of an application instance, uploads an HTML attack file, containing an attack script, to the repository of shared files of an application instance. The victim, another user of the same instance, downloads the file, causing the victim's browser to execute the script. The script may then be able to download other user files using an authentication cookie set in the victim's browser and send their contents to the attacker, who may not be authorized to see those other files. Furthermore, if the attack file is served with the same hostname as Web pages pertaining to the application instance (*application instance pages*) the same-origin policy may allow the attack script to read the contents of application instance pages that the victim but not the attacker is authorized to access. Worse, the script may thus be able to obtain a token embedded in a hidden field of an application instance form as a countermeasure against cross-site request forgery (CSRF) [9], and use it to defeat the countermeasure. The script can then submit the form, thus allowing the attacker to act upon the application instance while impersonating the victim, and with the privileges of the victim.

As a slight variation on the basic attack, the attacker may plant the attack file in advance of an anticipated revocation of the attacker's rights to access the application instance. The attacker may thus continue to read user files, read application instance pages, and interact with the application instance after revocation. If the victim's privileges include the right to create application instance accounts for new users, the attacker may be able to restore his or her lost access rights. This variation is of concern even for applications where all users have the same privileges and are authorized to access all files, a case in which cross-user attacks would seem at first glance to be inconsequential.

With a more elaborate variation, the attacker may be able to gain access to an application instance to which the attacker has not had access before. This requires some social engineering. The attacker wants access to a target application instance to which the victim has access. The attacker uses a second

application instance, to which the attacker has access, as a decoy; the decoy instance may already exist, or the attacker may register with the ASP and arrange for its creation. The attacker lures the victim into obtaining a user account for the decoy instance and logging in to the decoy instance while being logged in to the target instance. The attacker uploads an attack file with an attack script to the repository of the decoy instance and lures the victim into downloading the attack file. If the same hostname is used to serve user files and application pages pertaining to all instances, the same origin policy allows the attack script to access the target instance, read user files and application pages of the target instance and forward their contents to the attacker, interact with the target instance, and even, if the victim has sufficient privileges, create a user account at the target instance for him or herself.

4. Countermeasure

The same-origin policy prevents a script contained in an HTML document from accessing the contents of a Web page or file downloaded from a *different origin*. Assuming that all relevant pages or files are downloaded using the same protocol and port (typically *https* and 443 respectively), the *origin* of a page or file is, in effect, the DNS domain from which the page is served, or any domain that contains it as a subdomain. In other words, the origin is the hostname portion of the URL used to retrieve the page or file (henceforth, the *hostname of the page or file*), or a broader domain to which that hostname belongs.

The same-origin policy can be used to prevent attacks through HTML shared files if the hostnames of user files and application pages are chosen so that they satisfy the following two conditions:

- (1) The hostname of a user file of an application instance is different from, and is not a subdomain of, the hostname of any application page of the same or any other application instance. This ensures that an attack script in a user file will not be able to read the contents of any application page. It prevents an attacker, whether a current user or a former user, from using the script to obtain application data or to defeat CSRF countermeasures.
- (2) The hostname of a user file of an application instance is different from, and is not a subdomain of, the hostname of any user file of the same or any other application instance. This ensures that an attack script in a user file will not be able to read the contents of any other user file. It prevents the attacker from using the script to circumvent a system of file permissions or to continue reading user files after the attacker's credentials for accessing the application instance have been revoked.

These conditions could be satisfied, for example, by using user file hostnames of the form

userfiles-AppInstID-FileID.RegisteredDomain

and application page hostnames of the form

application-AppInstID.RegisteredDomain

where *AppInstID* is a unique application instance identifier, *FileID* is a unique file identifier, and *RegisteredDomain* is a domain name obtained by the ASP from an Internet domain registrar.

This hostname scheme, however, raises serious difficulties.

First, users may want to place links to an HTML user file in other HTML files (located anywhere in the World Wide Web). It is unreasonable to ask users to embed a file identifier in the hostname portion of the URL when constructing such a link. The file should instead be identified by an ordinary file path appended to a hostname that is the same for all user files pertaining to the same application instance.

Second, HTML user files belonging to the same application instance should be linkable to each other using relative URLs. This requires, again, that all user files pertaining to the same application instance have the same hostname.

Last but not least, an authentication cookie set upon login in a user's browser, if associated with the domain

`application-AppInstID.RegisteredDomain,`

cannot be used to authenticate user files that use different hostnames.

The solution to these difficulties lies in the combination of two ideas:

- (1) Use two different URLs to retrieve a user file, with a redirection step mapping the first one to the second.
- (2) Set two different authentication cookies upon login, one for user files and the other for application pages.

Of the two different user file URLs, the first one is a *standard URL* consisting of a *standard hostname* of the form

`userfiles-AppInstID.RegisteredDomain`

and an ordinary file path that identifies the file. Here is an example of a standard URL:

`https://userfiles-123456.pomcor.com/folder1/folder2/file1.html`

Standard URLs are reasonable ones to use in links to user files, and since all standard URLs for a given application instance have the same hostname, it is possible to use relative URLs in links to user files from other user files.

The second user file URL is an *extended URL* consisting of an *extended hostname* of the form

`userfiles-AppInstID-FileRetrSessID.RegisteredDomain`

and the same file path, e.g.:

`https://userfiles-123456-c4482d6f9e85d1a2c6441ea447c4167211ddc321.pomcor.com/folder1/folder2/file1.html`

where *FileRetrSessID* is a random, high-entropy, *file-retrieval session ID*. When the application receives a standard URL request, accompanied by a user file authentication cookie, it generates this ID and uses it as the primary key of an ephemeral file-retrieval session record that contains a reference to the login session identified by the cookie; then it redirects the browser to the extended URL. When the extended URL is received, the application uses the ID embedded in the hostname to authenticate the request by locating the file-retrieval session record and the login session record that it refers to; then it deletes the file-retrieval session record.

Besides being used for authentication, the file-retrieval session ID present in the hostname also serves to prevent an attack script contained in the user file from accessing any other user file, just like the above *FileID* would. Notice that it is not necessary for the ID to identify the file, or for the file-retrieval session record to contain a reference to the file. The file is identified by the file path component of the URL. The application verifies that the user has permission to download the file when the extended URL request is received, based on the file path and the user's user ID obtained from the login session record.

As for the two authentication cookies, their associated domains are the standard user file hostname and the application page hostname respectively. The application can set the user file cookie upon login by

redirecting the browser to a special URL whose hostname component is the standard user file hostname, and then redirecting it again to the URL of the initial application page. The special URL must have a path component that cannot be interpreted as the path of a user file; for example, user file and folder names beginning with a dot could be disallowed, and the path of the special URL could begin with a dot. The special URL must also contain an ephemeral secret, analogous to the ephemeral file-retrieval session ID, to authenticate the redirection.

5. Remarks

5.1 It is important that the extended user file hostname begin with the string "userfiles" or some other string that makes it clear that the hostname is used to serve user files, and the difference between user file hostnames and application page hostnames should be brought to the attention of the users. There are two reasons for this.

First, the extended user file hostname appears in the address box of the browser when an HTML user file is displayed. Making the hostname recognizable as a user file hostname will mitigate the risk of an HTML user file masquerading as the login page or some other application page.

Second, the warning issuing by the browser before running an executable user file will show the hostname of the file. Making the hostname recognizable as a user file hostname will mitigate the risk of an executable user file masquerading as a helper client-side program that the application is asking the user to install.

5.2 The countermeasure prevents a script contained in an HTML user file from reading any other user file. This imposes a constraint on the functionality of the private Web site comprised by the collection of HTML files of an application instance. It is possible to remove or relax this constraint at the cost of some reduction of security.

If there is no system of file permissions and all instance users are allowed to download all files, then the constraint could be removed by eliminating the extended user file hostnames and the user file redirection. The security cost would be the risk that a former user could retain access to user files.

If there is a system of file permissions based on a partition of the file space into several *security areas*, then the constraint could be relaxed by replacing the file-retrieval session ID with a security area ID. A script contained in an HTML user file could then access other files in the same security area. The security cost would be the risk that a former user could retain access to user files in the same security areas that user could access before his or her credentials were revoked.

5.3 The countermeasure remains useful in the case where there is only one application instance. (Of course, there are then no application instance IDs in hostnames.) There is no risk of cross-instance attacks in this case, but the countermeasure still protects against cross-user attacks and former user attacks.

6. References

- [1] [Multitenancy](http://en.wikipedia.org/wiki/Multitenancy). Wikipedia article, <http://en.wikipedia.org/wiki/Multitenancy>
- [2] [Cross-Site Scripting](http://en.wikipedia.org/wiki/Cross-site_scripting). Wikipedia article. http://en.wikipedia.org/wiki/Cross-site_scripting
- [3] [Same Origin Policy](http://en.wikipedia.org/wiki/Same_origin_policy). Wikipedia article. http://en.wikipedia.org/wiki/Same_origin_policy
- [4] [Risks of File-Sharing Technology](http://www.us-cert.gov/cas/tips/ST05-007.html). Cyber Security Tip ST05-007. <http://www.us-cert.gov/cas/tips/ST05-007.html>

- [5] MySecureCyberspace: [File Sharing](http://www.mysecurecyberspace.com/secure/file-sharing.html). <http://www.mysecurecyberspace.com/secure/file-sharing.html>
- [6] Web Security Mailing List. Thread: [File Uploading Vulnerabilities](http://www.webappsec.org/lists/websecurity/archive/2008-09/msg00022.html).
<http://www.webappsec.org/lists/websecurity/archive/2008-09/msg00022.html>
- [7] Gervase Markham, [Content Restrictions](http://www.gerv.net/security/content-restrictions/). <http://www.gerv.net/security/content-restrictions/>
- [8] Jesse Ruderman, [Security Tips for Web Developers](http://www.squarefree.com/securitytips/web-developers.html). <http://www.squarefree.com/securitytips/web-developers.html>
- [9] [Cross-Site Request Forgery](http://en.wikipedia.org/wiki/Cross-site_request_forgery). Wikipedia article. http://en.wikipedia.org/wiki/Cross-site_request_forgery