

Cross Site Request Forgery

An introduction to a common web application weakness

Jesse Burns

©2005, 2007, Information Security Partners, LLC.

<https://www.isecpartners.com>

Version 1.2

Introduction

Cross-site request forgery (CSRF; also known as XSRF or *hostile linking*) is a class of attack that affects web based applications with a predictable structure for invocation¹. This class of attack has in some form been known about and exploited since before the turn of the millennium. The CSRF name was given to them by Peter Watkins in a June 2001 posting to the Bugtraq mailing list.

CSRF flaws exist in web applications with a predictable action structure and which use cookies, browser authentication or client side certificates to authenticate users. The basic idea of CSRF is simple: an attacker tricks the user into performing an action of the attacker's choosing by directing the victim's actions on the target application with a link or other content. This is easiest to understand in the example of a HTTP GET.

For example, the link <http://www.google.com/search?q=iSEC+Partners> causes anyone who clicks it to search Google for "iSEC Partners". This is both harmless and by design. But a link like <http://www.isecpartners.com/EditProfile?action=set&key=emailAddress&value=evil@isecpartners.com> could tell an application which authenticated users only by cookie, browser authentication or certificate to edit a user's profile and change their email address.

Links can be easily obfuscated so they appear to go elsewhere, and to conceal words that would disclose their actual function. CSRF attacks effect applications that use either HTTP GET or HTTP POST to call their actions, although actions invoked with HTTP GET are often easier to exploit.

An Example of CSRF Exploitation — Goat Chat

Goat Chat is a hypothetical web application that offers messaging between users. Upon login Goat Chat sets a large, unpredictable session ID cookie which is used to authenticate further requests by users. One of the features of Goat Chat is that users can send each other links inside their text messages. Goat Chat is hypothetically deployed at <https://goatchat.isecpartners.com/> and uses HTTPS to keep messages, credentials, and session identifiers secret from network eavesdroppers. The application has effective cross-site scripting filters, blocking HTML content of any type.

Goat Chat supports the following user actions:

- Login
- Send a message
- Check for new messages
- Logout

An "incoming messages" frame is displayed to users who have logged in. This frame uses Javascript or a refresh tag to execute the "Check for new messages" action every 10 seconds on behalf of the user. New messages appear in this frame, and include the name of the sender and the text of the message. Text that is formatted as an HTTP or HTTPS URL is automatically converted into a link. The "Send a message" action takes two parameters: recipient (a user name or "all"), and the message itself which is a short string. The "Logout" action requires no parameters.

¹ Invocation means the calling of an action. An example action might be changing a setting, posting a search, or logging out of an application.

To determine if this application is susceptible to CSRF, we examine the “Send a message” action (although the “Logout” action is also sensitive and might be targeted by attackers for exploitation). When we do this, we find the following simple HTML form is submitted to send messages:

```
<form action="GoatChatMessageSender" method="GET">
Send To:<br>
<INPUT type="radio" name="Destination" value="Bob">Bob<BR>
<INPUT type="radio" name="Destination" value="Alice">Alice<BR>
<INPUT type="radio" name="Destination" value="Malory">Malory<BR>
<INPUT type="radio" name="Destination" value="All">All<BR>
Message: <input type="text" name="message" value="" />
<br><input type="submit" name="Send" value="Send Message" />
</form>
```

Figure 1 Form for sending a message

Here is what it looks like in its frame:



Figure 2 Rendering of Figure 1

From looking at this form we can figure out that when users wish to send the message “Hi Alice” to Alice, the following URL will be fetched when the user clicks Send Message.

```
https://goatchat.isecpartners.com/GoatChatMessageSender?Destination=Alice&message=Hi+Alice&Send=Send+Message
```

URL 1 An URL that sends “Hi Alice” to Alice when fetched by a logged in user

When the user performs an HTTP GET for URL 1, their browser includes the cookies appropriate for the goatchat.isecpartners.com site. These cookies are sent if the URL is typed in manually, if it is followed as part of loading a frame, clicking a link, due to an image request, or by submitting a form, even if it is loaded as the result of a 302 redirect or a meta-refresh tag. The only requirement for the cookie to be sent is that the logged in user is the one making a request. Attackers can exploit this.

If the logged in user were to visit a third party site run by devious hackers, and that site had the following image tag in it, URL 1 would be fetched by the victim’s browser and executed by the application server exactly as if the user submitted a message to Alice saying “Hi Alice”.

```
<img src =  
"https://goatchat.isecpartners.com/GoatChatMessageSender?Destination=Alice&message=Hi+Alice&Send=Send+Message" />
```

Figure 3 HTML which causes users to visit URL 1

An attacker could alternatively send the user a link in an email or via IM. The attacker could obfuscate it with a link to a site like <http://tinyurl.com/> which would then redirect the victim to the target site. Many sites have redirection pages that can be used to redirect victims to any other page, and the parameters are often subject to obfuscation via URL encoding or other tricks.

If this example had used the method POST rather than GET, then exploitation by link would be done differently. The attacker would send the victim a link that directed the victim to an attacker controlled site, and the site would contain (possibly within an iframe) a form pointing to the target site, but with hidden fields. The form could be submitted automatically with Javascript, or when the user clicked on it. More sophisticated variations allow for the exploitation even of multi-stage forms.

The example in Figure 4 demonstrates the exploitation of a system which allows password changes, but does not require the user's old password. In this example the target servlet requires an HTTP POST, and the attacker creates a self-submitting form to fulfill this requirement. Note that this exploit is not as reliable as the image based request in Figure 3 because the user's browser (or at least the tiny frame this exploit is placed in) is actually directed to the targeted site. Users with disabled browser scripting won't be exploited, and depending on the user's browser and configuration form submissions to other sites may result in a security popup box.

```
<HTML><BODY>  
<form method="POST" id="evil" name="evil"  
action="https://www.isecpartners.com/VictimApp/PasswordChange">  
<input type="hidden" name="newpass" value="badguy">  
</form>  
<script>document.evilm.submit()</script>  
</BODY></HTML>
```

Figure 4 Hidden form based exploit

Even if scripting is disabled however a "close this window" link that is actually a submit button may trick a user into submitting the form on the attacker's behalf.

Reflected vs. Stored CSRF

Similarly to Cross-site scripting (XSS) vulnerabilities, CSRF vulnerabilities can be divided into two major categories: *stored* and *reflected*.

A stored CSRF vulnerability is one where the attacker can use the application itself to provide the victim the exploit link or other content which directs the victim's browser back into the application, and causes attacker controlled actions to be executed as the victim. Stored CSRF vulnerabilities are more likely to succeed, since the user who receives the exploit content is almost certainly currently authenticated to perform actions. Stored CSRF vulnerabilities also have a more obvious trail, which may lead back to the attacker.

In a reflected CSRF vulnerability the attacker uses a system outside the application to expose the victim to the exploit link or content. This can be done using a blog, an email message, an instant message, a message board posting, or even a flyer posted in a public place with an URL that a victim types in. Reflected CSRF

attacks will frequently fail, as users may not be currently logged into the target system when the exploits are tried. The trail from a reflected CSRF attack may be under the control of the attacker, however, and could be deleted once the exploit was completed.

Easy Exploitation Scenarios

Some applications are easier to attack with CSRF than others. Certainly applications with stored CSRF vulnerabilities are reliably exploitable, but other decisions developers make also affect exposure.

Many applications direct HTTP GET calls to the same handler that is used for HTTP POST. In Java servlets, for example, the doGet() method simply calls the doPost() method, redirecting the parameters. This makes simpler image or link based exploits possible, and eases the exploitation of CSRF flaws.

Some applications, particularly intranet sites and administrative consoles in switches, access points, bridges, and other network devices use integrated browser authentication. This type of authentication doesn't expire, and the credentials remain available until the browser is closed. For someone who works with their browser all day, this can be hours or even days. This is a very wide window for attack, and this design decision increases the effectiveness of reflected CSRF attacks.

Many applications have extremely long cookie lives, allowing users to return to the site without re-authenticating. Long session lives can expose users to the risk of CSRF hours or even days after using a site. Popular sites with this configuration are of course even worse as attackers can guess that a large number of people selected at random are users of those services.

A few applications allow users to change their passwords without entering their old password. If these password change mechanisms are vulnerable to CSRF, then attackers may target this feature.

Distinguishing Between CSRF and XSS

CSRF and XSS (especially reflected XSS) are related security risks, and the similarity can be confusing. Many people find themselves trying to determine if an attack they have uncovered is exploiting an XSS weakness or an CSRF weakness. The distinction is often easier to make by considering the solution to the weakness you have identified. In the case of an XSS flaw, an attacker exploits a lack of input and / or output filtering. If a change to the application that filters out dangerous characters like <, >, ", ', &, ;, or # could resolve the flaw, then it is not an CSRF issue but an XSS issue. CSRF is about the predictability of the structure of the application. XSS is related to the application performing insufficient data validation.

XSS flaws may allow bypassing of any CSRF protections by leaking valid values of the tokens, allowing Referer headers to appear to be the application itself, or by hosting hostile HTML and Javascript elements right in the target application. Therefore resolving XSS flaws should be given priority over CSRF weaknesses – although CSRF is usually a fatal flaw as well.

Myths About CSRF

Myth: CSRF is just a special case of XSS.

Fact: CSRF is a separate vulnerability from XSS, with a different solution. XSS protections won't stop CSRF attacks, although XSS are important to solve and should be prioritized.

Myth: Applications aren't vulnerable to CSRF if they use multi-page forms to perform actions.
Fact: While multi-page forms certainly make exploitation harder, attackers can usually exploit them. iSEC frequently uses multiple iframes when demonstrating multi-page form CSRF exploits for customers.

Myth: Applications that use HTTP POST aren't vulnerable to CSRF.
Fact: While POSTs can be more difficult to exploit, they certainly are exploitable. Forms can mislead users about what they are sending and to where, and scripting can lead to automatic submission.

Myth: CSRF attacks are the user's fault.
Fact: Attackers can usually exploit CSRF vulnerabilities without the knowledge or consent of the user.

Myth: CSRF can be prevented by filtering based on the Referer header.
Fact: This is very unreliable since attackers can easily block the sending of the Referer header, and the HTTP RFCs make it clear the this header is optional. Browsers also omit the Referer header when making requests over SSL.

Myth: CSRF weaknesses are low risk.
Fact: CSRF weaknesses can be used by attackers to perform any action an authorized user of the application can perform! This could include changing passwords, buying merchandise, or transferring money, depending on the domain of your application.

Myth: My firewall / SSL server / or the .NET / Struts framework protect me from CSRF.
Fact: Firewalls and SSL provide security guarantees completely unrelated to ensuring that action requests were formulated by the application and submitted on purpose by the user. The .NET and Struts application frameworks make no reliable, documented guarantees of protection against CSRF. If you haven't specifically addressed the need to protect your application about CSRF, you are almost certainly vulnerable.

Terminology Definitions

- Action: Some behavior in the web application such as setting an account parameter, executing a trade, creating, updating, or deleting an object. An action is associated with a HTTP request that is usually a GET or POST, and some number of query parameters.
- Action Formulator: The source of the HTML elements that resulted in the creation of the request. Under normal circumstances a web application is the only action formulator for its actions, but in the case of a CSRF attack, an unauthorized action formulator is present.
- Action Name: A unique name for each action. This could be a human readable name like "Execute_Trade" or a unique number like 42. Action names can be implied from the URI they are located at; for example, the name of the servlet that implements the action could be its effective action name.
- Session Identifier: An unpredictable value used to uniquely identify a user's session. In a J2EE environment this is typically stored in the *JSESSIONID* cookie, while in an ASP environment this cookie is commonly called *ASPSESSIONID*. Many other web application frameworks have an unpredictable session identifier stored in a cookie, and is used to identify user sessions and store server side information about the user.
- HMAC_sha1(x, y): A keyed cryptographic hash based on SHA-1² and using the keying technique described in RFC 2104³. The output of this function needs to be encoded so that it will not contain

² <http://www.itl.nist.gov/fipspubs/fip180-1.htm> Secure Hash Algorithm 1. While recent cryptanalysis results suggest this algorithm is theoretically imperfect, it is the current standard for secure cryptographic hashes.

characters that are invalid in HTTP headers, for example with base 64 encoding⁴. The Java cryptography extensions⁵ and the OpenSSL⁶ and RSA BSAFE⁷ libraries all support HMACs with SHA-1.

- **Query Parameter:** a name/value pair, associated with a GET or POST request. Query parameters can contain values provided by the user, or values from the server generating the form or link the query was generated by.

Protection Approaches

Approach 1: Use cryptographic tokens to prove the action formulator knows a session- and action-specific secret.

Level of protection: Very High

Recommended by iSEC

Description: When the web application formulates an action (by generating a link or form that causes an action when requested by the user), it includes as a query parameter (usually as a hidden input tag) a name/value pair with a name like *CSRFPreventionToken* and the value $\text{HMAC_sha1}(\text{action_name} + \text{secret}, \text{session_id})$.

When the application receives an action request, but before executing the action, it verifies the value of *CSRFPreventionToken* by comparing the value of the provided token to a calculation of $\text{HMAC_sha1}(\text{requested_action_name} + \text{secret}, \text{session_id})$. If the values do not match, then the action formulator was not the application. The action should be aborted and the event can be logged as a potential security incident.

Note that *action_name* should be different for each action, although *secret* can remain constant. From time to time tokens may leak out from the application; for example, an application that doesn't use SSL and which stores the token in the query string will divulge the tokens to other sites in the Referer field when users click on links to those sites. By keeping *action_names* unique, the application minimizes the impact on the application of a potentially hostile third party learning *CSRFPreventionToken* value (the attacker could now formulate only the actions with the same name as that which referred the victim to the attacker's site). Using SSL is almost always necessary for secure web applications.

Advantages: Very strong protection, no additional memory requirements per user session.

Disadvantages: Requires the dynamic generation of all actions. This widespread change can be eased through integration with a thin client framework. The approach also requires a small amount of computation when actions are formulated and verified.

³ <http://www.faqs.org/rfcs/rfc2104.html> Keyed-Hashing for Message Authentication. This RFC defines a standard way of creating cryptographic hashes that can only be verified with knowledge of the key.

⁴ <http://en.wikipedia.org/wiki/Base64> A common web application encoding format for binary data.

⁵ <http://java.sun.com/products/jce/> The standard for java crypto, with many provider implementations including one from Sun.

⁶ <http://www.openssl.org/> A common, free cryptography library. This library has had numerous security flaws, and should be kept up to date.

⁷ <http://www.rsasecurity.com/node.asp?id=1202> A commercial, somewhat expensive cryptography provider for Java and C.

Approach 2: Use secret tokens to prove the action formulator knew an action- and session-specific secret.
Level of protection: Very High **Recommended by iSEC**

Description: When the web application formulates an action, it includes as a query parameter (usually as a hidden input tag) a name/value pair with a name like *CSRFPreventionToken*, and a value that is a high-entropy 128 bit value that has been base 64 encoded. For each action/session pair this token need only be generated once, and the token value is stored in a session-specific table mapping action names to validation tokens.

When the application receives an action request, but before executing the action, it verifies the value of *CSRFPreventionToken* by comparing the provided token to the value stored in the mapping table for this action. If there is no value in the table for this action name, or if the value provided does not match the value in the table exactly, then the action formulator was not the application. The action should be aborted and the event can be logged as a potential security incident.

Similarly to Approach 1, the action name should be different for each action, and for the same reason.

Advantages: Very strong protection, minimal computational overhead.

Disadvantages: Requires the dynamic generation of all actions. This widespread change can be eased through integration with a thin client framework. Requires additional memory on the order of 128 bits times the number of actions per session.

Approach 3: Use the optional HTTP Referer header⁸ to verify action formulators.
Level of protection: Medium

Description: When the application receives an action request, but before executing the action, it verifies that the value of the Referer header in the HTTP request is from a source that is authorized to invoke the application. Disallow access if the Referer header is not present, or if its value points to a page or site that should not be formulating the action. If the header isn't from the allowed set, the action should be aborted and the event can be logged as a potential security incident.

Advantages: Simple to implement.

Disadvantages: The header is optional and may not be present – some browsers disable this header, and it is not available when interactions occur between HTTPS and HTTP served pages. The risk of header spoofing exists, and tracking the valid sources of invocations may be difficult in some applications.

⁸ <http://www.w3.org/Protocols/HTTP/HTTRQ-Headers.html#z14> this header was misspelled referer in the HTTP\1.1 specification, and now must be given with the spelling seen here.

Approach 4: Require changes to application state to be done only with HTTP POST operations.

Level of protection: Very Low

Description: Only allow HTTP POST operations to perform changes to the state of the application such as creating, updating, or deleting of objects. This makes exploitation through image tags ineffective, and helps reduce exploitability of allowing image links to be shared inside an application. In a low risk environment, this may encourage attackers to target other sites rather than those with this limited protection.

Advantages: Simple to implement.

Disadvantages: Attackers can adjust their attacks to be form-based like CSRF examples 1 and 2, submitting forms automatically or tricking users by making huge, mislabeled submit buttons.

Approach 5: Use a simplified `CSRFPreventionToken`.

Level of protection: Medium to High

Description: This is just like Approach 1, except that a single `CSRFPreventionToken` is generated when the user logs in, is stored in the session state, and is valid for *all* actions performed in the context of that session (but not in the context of other sessions, and not when the session expires).

Whenever the application receives an action request, it checks that the provided token matches the correct token in the session state. As with the other approaches, it aborts the action if the tokens do not match.

As always, the token must be kept secret (by e.g. not sending it to the server in the query string).

Advantages: Simpler to implement than Approaches 1 or 2.

Disadvantages: Requires changes to all action formulations in the application, requires the application to be completely under SSL, or to use HTTP POST only for any operation that required a `CSRFPreventionToken` in its formulation to avoid token leakage through the Referer header.