1.5.5 Cookie Analysis

As part of the first audit, Facebook were asked to provide an explanation of the purpose of each of the identified cookies. The information provided at the time has been reviewed and updated as part of this audit. It was noted at the time of the first audit that Facebook uses many cookies for many purposes and it is not feasible to identify and analyse the purpose of every single cookie. Therefore, the focus of the analysis continues to be on the cookies identified in the previous sections.

Some of the cookies used by Facebook are known as session cookies. In the majority of cases, these cookies remain on the user's PC until the web browser is exited. There are a few scenarios, as mentioned in the first report, such as Firefox session restore mode where session cookies can be retained after the browser has been exited.

1.5.5.1 datr

The purpose of the datr cookie is to identify the web browser being used to connect to Facebook independently of the logged in user. This cookie plays a key role in Facebook's security and site integrity features.

At the time of the first audit, the datr cookie generation code was reviewed and it was confirmed that the execution path followed in the case of a request for social plugin content does not set the datr cookie.

The lifetime of the datr cookie is two years.

1.5.5.2 reg_fb_gate, reg_fb_ref and reg_fb_ext

The reg_fb_gate cookie contains the first Facebook page that the web browser visited. The reg_fb_ref cookie contains the last Facebook page that the web browser visited. The reg_fb_ext cookie contains an external referrer URL form when the browser first visited Facebook.

These cookies are only set when the browser is either not a Facebook user or is not logged in to Facebook. These cookies are used by Facebook to track registration effectiveness by recording how the user originally came to Facebook when they created their account.

These three cookies are session cookies.

1.5.5.3 wd

This cookie stores the browser window dimensions and is used by Facebook to optimise the rendering of the page.

The wd cookie is a session cookie.

1.5.5.4 c_user

This cookie contains the user ID of the currently logged in user.

The lifetime of this cookie is dependent on the status of the 'keep me logged in' checkbox. If the 'keep me logged in' checkbox is set, the cookie expires after 30 days of inactivity. If the 'keep me logged in' checkbox is not set, the cookie is a session cookie and will therefore be cleared when the browser exits.

1.5.5.5 lu

The lu cookie is used to manage how the login page is presented to the user. Several pieces of information are encoded within the lu cookie.

The 'keep me logged in' checkbox on the Facebook login page is used to determine whether or not the authentication cookies delivered to the user when they log in will be retained when the user quits their browser. If the 'keep me logged in' checkbox is ticked, then when the user logs in the authentication cookies will be persistent (retained after the browser exits). If the 'keep me logged in' checkbox is not ticked then the authentication cookies will be session cookies (cleared when the browser exits).

The user can explicitly check or uncheck the 'keep me logged in' box. The lu cookie records whether the user has performed such an explicit action.

If the user has not explicitly either checked or unchecked the 'keep me logged in' box, then the default mode of operation is to automatically check the 'keep me logged in' box if the same user has logged in from the same computer three times in a row without logging out. A user explicitly checking or unchecking the 'keep me logged in' box always overrides this feature.

To implement this functionality, the lu cookie contains a counter which is incremented if the user logging in is the same as the previous user that logged in from this web browser, and if the previous user did not explicitly log out. To be able to determine whether the user logging in is the same as the previous user that logged in, the lu cookie contains the user ID of the previously logged in user. The previously logged in user component of the lu cookie is set to zero if the user explicitly logs out.

The user ID component of the lu cookie is also used to pre-populate the email address field of the login form if the user did not previously explicitly log out.

To summarise, the components of the lu cookie are:

- The user ID of the previously logged in user, or zero if the user explicitly logged out.
- A counter containing the number of times in a row that the same user has logged in from this browser and has not explicitly logged out.

- A flag to indicate whether the user has explicitly either checked or unchecked the 'keep me logged in' box.

The lifetime of the lu cookie is two years.

1.5.5.6 sct

At the time of the first audit the presence of a cookie named sct was noted. This cookie contained a unix timestamp value representing the time at which the user logged in. This cookie was used to distinguish between two sessions for the same user, created at different times.

The absence of this cookie was noted at the time of the second audit and it has been confirmed by Facebook that the unix timestamp value previously contained in the sct cookie has been incorporated into the xs cookie described in the next section.

1.5.5.7 xs

This cookie contains multiple pieces of information, separated by a colon[19].

At the time of the first audit it was noted that the values contained within the xs cookie were;

- The first portion is an up-to-two digit number representing the session number.
- The second portion is a session secret.
- The third, optional, portion is a secure flag, which is used if the user has enabled the secure browsing feature.

It was noted at the time of the second audit that the xs cookie now contains four components separated by colons. The first three components are consistent with the three functions described above and the fourth component appears to be a unix timestamp, consistent with the incorporation of the value previously carried by the sct cookie into the xs cookie.

1.5.5.8 x-referer

This cookie contains the full referrer URL.

When a user clicks on a link on a web page, this leads to a HTTP request being sent to a server. The referrer is the URL of the web page on which the link that the user clicked resided. The referrer is sent with every HTTP request[20].

Facebook use this value to correctly capture the referrer for pages using Facebook Quickling navigation. Quickling navigation is a feature that uses AJAX to make Facebook page requests,

[19] Colon is encoded to the value %3A for transmission.

[20] http://tools.ietf.org/html/rfc2616#section-14.36

thus speeding up the user experience of the site[21]. In these cases, the actual referrer URL is in the URL fragment[22] and this is not normally sent to the server in the HTTP Referer[23] header.

1.5.5.9 presence

The presence cookie is used to contain the user's chat state. For example, which chat tabs are open. This is a session cookie.

1.5.5.10 p

The p cookie is known as the user's channel partition and is required by many features on the Facebook site, including chat and client-side notifications. This is a session cookie.

1.5.5.11 locale

This cookie contains the display locale of the last logged in user on this browser. This cookie appears to only be set after the user logs out and has a lifetime of one week.

1.5.5.12 lsd[24]

At the time of the first audit it was reported that the lsd cookie contains a random value that is set when a Facebook user logs out to prevent cross-site request forgery (CSRF) attacks.

Cross-site request forgery is an attack technique involving misuse of credentials from one site (in this case Facebook) to perform unauthorised actions on a user's account when the user visits a web site containing specifically crafted malicious code.

Further insight into the operation of this cookie has been gained on when this cookie is set as part of this audit. In particular, the cookie is not just set when a Facebook user logs out, rather, the cookie is set whenever the browser is in a logged out state.

The lsd cookie is a session cookie.

---

[21] Some technical detail can be found at http://www.slideshare.net/ajaxexperience2009/chanhao-jiang-and-david-wei-presentation-quickling-pagecache

[22] The URL fragment is the name given to the part of the URL after a "#" and is typically, but not always, used to refer to a part or position within a HTML document. See http://tools.ietf.org/html/rfc3986

[23] The HTTP referrer header is mis-spelled as "Referer" in the HTTP standard, so this is the correct name of the HTTP header as per the standard.

[24] FB-I reports that in the period between the time this cookie testing was performed and the completion of the report, the lsd cookie has been removed.

1.5.5.13 Cookies beginning with _e_

At the time of the first audit, it was noted that a substantial number of cookies that begin with the characters "_e_" were transmitted. These were referred to by Facebook as EagleEye cookies.

The cookie names consisted of "_e_" followed by a four character random string, followed by an underscore and then an incrementally increasing number, starting at zero. For example, _e_gh2c_0, _e_gh2c_1, _e_gh2c_2, etc.

It was reported that these cookies were generated by JavaScript and used to transmit information to Facebook about the responsiveness of the site for the user. Cookies were being used as the transport mechanism for the performance related information, but the content of the cookies was being generated in the user's web browser and no information was being transferred to Facebook that was not available for transmission in some other form (e.g. in a HTTP POST). Facebook did not place any information on the user's PC using these cookies.

It was further noted that it was possible to observe, by monitoring the communication between the web browser and Facebook, that each time an EagleEye cookie was submitted to Facebook, the corresponding response unset that cookie. This is consistent with the explanation provided by Facebook that these cookies are used as a transport mechanism.

The EagleEye cookie value consisted on an encoded JSON structure that contained information about an action performed by the user on the site. For example, when the user clicked on a link.

The testing carried out as part of this audit revealed that _e_ cookies are still in use and their behaviour continues to be as described here.

1.5.5.14 fr

As part of the testing carried out for this audit, a new cookie named fr was identified.

It was noted that the cookie is only set when a Facebook user logs in to the site and it has an expiry period of 30 days. The cookie, including the encrypted user id, is retained after the user logs out. Upon examination, the fr cookie clearly consists of two components.

An example fr cookie value is "0nx07ppspaoOQlQv1.AWVAlyAiGNl9vuExmcrX2lmfAfk".

Facebook were asked to provide an explanation of the purpose of this new cookie.

The content of the two parts of the cookie have been reported to be as follows;

- The first part of the cookie is a browser ID, used to identify the web browser.
- The second part of the cookie is an encrypted version of the logged in user's Facebook ID. The user's ID is re-encrypted every hour to a different value.

The code used to generate the fr cookie value has been reviewed and it has been confirmed that the browser ID is a random value and the encrypted user ID value contains only the Facebook user ID. It was also confirmed by code review that the fr cookie value generation code is called whenever the other login session cookie values (c_user, xs, etc.) are refreshed, which takes place roughly hourly but this can vary for operational reasons.

This cookie is being used by Facebook to deliver a series of new advertisement products such as real time bidding, which works as follows:

- An advertising partner of Facebook, for example, doubleclick has an ad on, for example, the New York Times website[25].
- A Facebook user visits the New York Times.
- The website contains a pixel image which causes a request to be sent to Facebook. Usually, the request to Facebook will have a referrer value of the partner (in this case doubleclick), along with an opaque partner value provided by doubleclick. In some cases, partners do not control the browser referrer value. In such cases, FB-I states that they exclude this referrer value from their impression logs and do not use it as part of this or other advertising systems.
- Facebook store a relationship between the partner value and the fr cookie browser component value.
- Then, when the user visits Facebook, the partner is sent the partner value and can respond with a bid amount to bid to have an ad displayed to the user.
- If the partner wins the bid, Facebook will serve a standard ad from a standard ad campaign to the user.

To summarise what each of the actors know about the user's activity:

- The partner (doubleclick in this case) knows that the user has visited the New York Times website.
- Facebook do not know that the user has visited the New York Times website[26]. The meaning of the partner value provided to Facebook is opaque to Facebook. FB-I report that the partner values are typically short identifiers. While this value may, hypothetically, somehow encode the fact that the user has visited the New York Times website it is not clear how or why the partner would choose to do this. The partner will store whatever data they know about the user in their own database.

[25] With the exception of Facebook, the actors described in the following example are intended purely to illustrate the functionality of the cookie. It is not known, nor is it relevant, whether doubleclick is an advertising partner of Facebook or whether doubleclick have an ad on the New York Times website.

[26] As mentioned above, sometimes FB-I may receive this value in a HTTP referrer header but they have stated that they do not log the referrer value from such requests.

- Facebook know information about the user provided separately by the user to Facebook (e.g. the user's profile information).
- The partner has no access to any information provided by the Facebook user to Facebook.
- Due to the bid requests, the partner may know which browsers are active Facebook users, but they are contractually prohibited from storing or using this fact.

Although this cookie will be sent in requests for social plugins that occur after a browser has had a logged in user, FB-I states that this cookie is not currently used other than as described above.

1.5.5.15 sub

During the testing for this audit, the presence of a cookie named sub was noted. This cookie was not present at the time of the first audit. The value of the sub cookie was noted to be a simple numeric value but Facebook were asked to clarify the purpose of this new cookie.

The chat functionality on the Facebook site works using a technique known as HTTP long polling[27]. This technique involves the client sending a HTTP request to the chat server and the server holding the connection open by taking a long time to respond to the request.

This leads to a situation where, if the user has multiple tabs open, there are multiple simultaneously open HTTP connections to the same server. Most browsers limit the number of allowed simultaneous connections (typically to a value somewhere in the region of six).

The sub cookie is used by Facebook's chat JavaScript to communicate across tabs to coordinate connections to the Facebook chat server. The sub cookie replaces an older, less effective technique for addressing the same issue.

1.5.6 Active Cookie Management

As part of the first audit, Facebook demonstrated a feature for proactive management of browser cookie state, known as "Cookie Monster".

The cookie management framework contains configuration for each cookie and the context in which the cookie should be set. For example, certain cookies are required in the context of a logged in user and after the user logs out these cookies should be unset. If a cookie is received for which there is not a configuration, it will automatically be cleared.

The cookie management framework is executed on every Facebook request, including requests from social plugins. Unexpected cookies, or cookies from the incorrect context (such as cookies that are only meaningful in the context of a logged in user being received in a request from a non-logged in user), are automatically unset.

[27] http://en.wikipedia.org/wiki/Push_technology#Long_polling