The first HTTP request to Facebook for social plugin content transmitted the three cookies listed above, along with a cookie named wd, presumably generated by JavaScript. The wd cookie has the value “1082x676” which represents the dimensions of the browser window in which the Facebook page was loaded. The HTTP response received from Facebook in response to this first HTTP request unsets the wd cookie.

The three cookies listed above are transmitted in each of the remaining 11 captured HTTP requests to www.facebook.com for social plugin content.

This finding was confirmed by examining the content of the folder C:\Documents and Settings\<User>\Cookies (where <User> is the username of the currently logged in Windows user). Since the reg_fb_gate and reg_fb_ref cookies are session cookies, it would not be expected that they would be found in the Cookies folder. Indeed, only an entry for the datr cookie is found in the Cookies folder.

In this case, where a non-Facebook user visits www.facebook.com, without registering or logging on, four cookies are set. Explanations of the purpose of these cookies can be found in the cookie analysis section below.

6.4 Facebook Users and Cookies

Using a similar technique to the one described above, the cookies sent to Facebook when a logged in or logged out Facebook user browses to sites containing social plugins have been identified.

The testing was performed on a newly installed, fully patched Windows XP virtual machine with anti-virus software installed. All browsing was carried out using the default configuration of Internet Explorer 8 (Version: 8.0.6001.18702).

All traffic generated by the virtual machine was captured by an instance of Wireshark running in the host operating system. The captured packet data was examined and the HTTP requests/responses associated with retrieving social plugin content from Facebook were examined. The cookies sent by the web browser were identified and are described in the cookie analysis (Section 6.5).

6.4.1 Logged In Users

The Facebook website was visited and a user account was used to log in. A period of browsing activity to non-Facebook sites, some with social plugins and some without, then took place. Each request to Facebook for social plugin content transmitted the same set of cookies:

- datr
- c_user
- lu
- sct
- xs
- x-referer
- presence
- p

The purpose of each of these cookies is discussed in Section 6.5.

6.4.2 Logged Out Users

The Facebook website was visited again and a user account was logged out. A period of browsing activity to non-Facebook sites, some with social plugins and some without, then took place. Each request to Facebook for social plugin content transmitted the same set of cookies:

- datr
- lu
- x-referer
- locale
- lsd
- reg_fb_gate
- reg_fb_reg

The purpose of each of these cookies is discussed in Section 6.5.

6.5 Cookie Analysis

Facebook have been asked to provide an explanation of the purpose of each of the cookies identified. The information provided below is correct at the time of writing but is subject to change over time. Facebook uses many cookies for many purposes, and it is not feasible as part of this report to identify and analyse the purpose of every single cookie. Therefore, the focus of the following analysis is on the cookies identified in the preceding sections.

The lifetimes of each of the cookies is provided below.

Some of the cookies in the following sections are referred to 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 such as Firefox session restore mode where session cookies may be retained after the browser has been exited[24].

6.5.1 datr

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

The datr cookie generation and setting code has been reviewed and it has been 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 currently two years.

6.5.2 reg_fb_gate, reg_fb_ref and reg_ext_ref Cookies

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.

[24] See http://support.mozilla.com/en-US/kb/Session%20Restore#w_when-session-restore-occurs for more details.

As described above, these cookies appear to only be 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.

The functionality of these cookies has been verified as follows:

- Using a newly installed, fully patched Windows XP virtual machine with anti-virus software installed. All browsing was carried out using the default configuration of Internet Explorer 8 (Version: 8.0.6001.18702).
- All traffic generated by the virtual machine was captured by an instance of Wireshark running in the host operating system.
- The URL “http://www.facebook.com/imdb” was typed into the browser.
- The URL “http://www.facebook.com/VultureCentral” was typed into the browser.
- The captured packet data was examined and the HTTP requests/responses associated with the two requests above were identified. It was noted that:
    - The response to the HTTP request for /imdb sets the reg_fb_gate and reg_fb_ref cookies as follows:
        - reg_fb_gate = http%3A%2F%2Fwww.facebook.com%2Fimdb
        - reg_fb_ref = http%3A%2F%2Fwww.facebook.com%2Fimdb
    - The response to the HTTP request for /VultureCentral further sets the reg_fb_ref cookies as follows:
        - reg_fb_ref = http%3A%2F%2Fwww.faceboom.com%2FVultureCentral

The reg_ext_ref cookie value contains an external referrer URL from when the browser first visited Facebook. The functionality of this cookie has been verified as follows:

- Using a newly installed, fully patched Windows XP virtual machine with anti-virus software installed. All browsing was carried out using the default configuration of Internet Explorer 8 (Version: 8.0.6001.18702).
- All traffic generated by the virtual machine was captured by an instance of Wireshark running in the host operating system.
- The URL “http://www.google.com” was entered into the browser.
- The search term “Guinness _acebook” was entered.
- The search result for the Guinness Facebook page was clicked.
- The captured packet data was examined and the HTTP requests/responses associated with the HTTP request for the Guinness Facebook page was identified. It was noted that
    - The response to the HTTP request for /GuinnessWorldRecords sets the reg_ext_ref cookie to a Google URL.
    - The reg_fb_gate and reg_fb_ref cookies are set consistent with the testing described above.

The reg_fb_gate, reg_fb_ref and reg_ext_ref cookies are session cookies.

6.5.3 The wd Cookie

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

The functionality of this cookie has been verified as follows:

- Using a newly installed, fully patched Windows XP virtual machine with anti-virus software installed. All browsing was carried out using the default configuration of Internet Explorer 8 (Version: 8.0.6001.18702).
- All traffic generated by the virtual machine was captured by an instance of Wireshark running in the host operating system.
- Visit “http://www.facebook.com”
- Reload Facebook web page
- Note that HTTP request for Facebook page sends cookie wd with value “771x404”
- Make browser window larger
- Reload Facebook web page
- Noted that the HTTP request for the Facebook page sends cookie wd with value “953x453”

Used website http://whatsmy.browsersize.com/ to verify that the values provided in the wd cookie are consistent with the browser window dimensions. The window dimensions reported by browsersize.com are consistently 21 pixels larger than those contained in the wd cookie. The reason for this discrepancy has not been investigated, but it is not considered important for the purpose of the current testing. It is believed to have been adequately demonstrated that the wd cookie represents the window dimensions.

The wd cookie is a session cookie.

6.5.4 c_user

The c_user 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.

6.5.5 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, as described here.

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.

In order 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[25]. 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 logs out.
- A counter containing the number of times in a row that the same user has logged in from from this browser and has not explicitly logged out.
- A flag used 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.

6.5.6 sct

The sct cookie contains a unix timestamp value[26] representing the time at which the user logged in. This cookie is used to distinguish between two sessions for the same user, created at different times.

The value contained in the sct cookie has been verified to be consistent with the time and date at which test logins were performed.

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.

6.5.7 xs

This cookie contains multiple pieces of information, separated by colon[27]. The first value is an up to two-digit number representing the session number. The second portion of the value is a session

[25] For example, if the user closed their browser rather than explicitly logged out.

[26] The value is defined as the number of seconds elapsed since midnight UTC of January 1, 1970, not counting leap seconds.

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

secret. The third, optional component is a 'secure' flag for if the user has enabled the secure browsing feature.

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.

6.5.8 x-referer

This cookie contains the full referrer URL[28]. Facebook use this value to correctly capture the referrer for pages using Facebook Quickling navigation[29]. In these cases the actual URL is in the URL fragment[30] and this is normally not sent to the server in the HTTP Referer[31] header.

6.5.9 presence

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

This cookie is a session cookie.

6.5.10 p

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

This cookie is a session cookie.

6.5.11 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.

The locale cookie has a lifetime of one week.

6.5.12 lsd

The lsd cookie contains a random value that is set when a Facebook user logs out in order to prevent cross-site request forgery (CSRF) attacks.

[28] 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 on resided. The referrer is sent with every HTTP request. See http://tools.ietf.org/html/rfc2616#section-14.36.

[29] Quickling navigation is a feature that uses AJAX to make Facebook page requests to speed up the user experience of the site. Some technical detail can be found here: http://www.slideshare.net/ajaxexperience2009/chanhao-jiang-and-david-wei-presentation-quickling-pagecache.

[30] 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.

[31] The HTTP referrer header is misspelled as "Referer" in the HTTP standard, so this is the correct name of the HTTP header as per the HTTP standard.

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

The lsd cookie is a session cookie.

6.5.13 Cookies Beginning with _e_

When monitoring the communication between Facebook and a web browser it is possible to note that a substantial number of cookies that begin with the characters "_e_" are transmitted. These are referred to by Facebook as EagleEye cookies.

The cookie names consist 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.

These cookies are generated by Javascript and used to transmit information to Facebook about the responsiveness of the site for the user.

Cookies are being used as a transport mechanism for the performance related information, but the content of these cookies is all being generated by the user's web browser and there is no information being transferred to Facebook that is not available for transmission in some other form (e.g. in a HTTP POST). Facebook do not place any information on the user's PC using these cookies.

It is possible to observe, by monitoring the communication between the web browser and Facebook, that each time an EagleEye cookie is submitted to Facebook, the corresponding response will unset that cookie[32]. Since the cookie is only serving as a transport mechanism to deliver the performance related information to Facebook, once the cookie has been successfully received by Facebook it serves no further purpose and can be deleted.

The EagleEye cookie consist of an encoded JSON[33] structure that contains information about an action performed by the user on the site. For example, when the user clicks on a link.

6.6 Active Cookie Management

Facebook have demonstrated a recently improved feature for proactive management of browser cookie state, known as "Cookie Monster". The code of this feature has been reviewed and confirmed to operate as described in this section.

Historically, the deletion of cookies on logout required manual insertion of code into the logout process to unset each cookie. This technique was error prone, since developers could add a new cookie but forget to add the corresponding code to unset the cookie on logout.

[32] It is possible under some circumstances, as described in Section 0, that the HTTP response is delivered to the server before the cookie management code is executed. In these cases, the EagleEye cookies will be deleted the next time the cookie management code is executed for this user.

[33] http://tools.ietf.org/html/rfc4627