Appendix 1

Facebook Technical Analysis Report

Prepared for the Data Protection Commissioner
By Dave O'Reilly

16th December 2011

1. Introduction

The purpose of this document is to provide a technical analysis of certain aspects of Facebook's architecture, infrastructure and functionality. The focus of this report is on how the various features that were studied operate.

Wherever possible, sources of evidence have been sought and experiments carried out to validate that the features perform as described.

Every effort has been made to make the test results produced in this report as repeatable as possible.

Unless otherwise mentioned, all tests were performed in 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). A snapshot of the newly installed virtual machine state was taken and the snapshot was restored before each test described within this document, except where explicitly explained otherwise.

New test Facebook accounts were created as required.

In order to verify certain claims, aspects of the Facebook source code have been examined. Source code examination took place by examining the content of the Facebook source code repository. All examinations were carried out on the trunk of the repository, representing the currently deployed code base.

The code examined was PHP, which is compiled into C++ binaries for deployment.

2. Contact Importing

2.1 Background

When a user creates a Facebook account, they have the opportunity to import their contacts from a range of email service providers into Facebook. It is possible that the user's contacts will include both users and non-users of Facebook. As well as sending friend requests to existing Facebook users, the user performing the contact import has the opportunity to invite the non-users to join Facebook and become friends.

If the user sends an invitation to a non-user, this will cause the non-user to receive an email from Facebook containing a link that will allow the non-user to create a Facebook account.

The non-user can ignore this email if they do not want to join Facebook. A link is provided in the invitation email that allows the non-user to choose to opt out of receiving subsequent invitation requests from Facebook.

It is possible that a second Facebook user could import the same non-user email address. Assuming that the non-user does not choose to opt out of receiving invitations, a second invitation could be sent to the non-user by the second Facebook user. The second (and subsequent) invitations may include reference to other Facebook users that the non-user may know.

2.2 Storage and Removal of Contact Data

The data structures within which imported contact information is stored have been reviewed. No distinction was apparent in the storage structures based on whether the contact information was that of an existing Facebook user or a non-Facebook user. These data structures are not the same as the data structures used to store a Facebook user's profile.

The imported contact information appears to be stored in the following way(s):

- Each time a user performs an import, the imported data is added to an array of imports, one entry for each set of imported data. Each entry in this array consists of a data structure containing an array of the contact names and a corresponding array of the contact email addresses. This information is associated with the importing user's Facebook account.
- A data structure consisting of a hash of the email address of the imported contact and a string consisting of a comma separated list of Facebook user IDs for users that have imported that particular email address.
- The contact information is stored in the data structure representing the user's address book.
- The contact information is also stored in the data structure representing the user's phone book.

No other techniques for the storage of contact information about non-Facebook users have been identified.

By examining the source code it has been confirmed that upon receipt of a user request to remove all imported contact data, the following steps are carried out:

- The data is removed from the array of imports.
- The Facebook user ID of the user requesting the removal of the imported data is removed from the comma separated list of user IDs associated with all of their contact email addresses. If there are no remaining user IDs associated with a particular contact email address, the contact email address entry is also removed[1].
- All imported contact data is removed from the user's address book.
- All imported contact data is removed from the user's phone book.

2.3 Use of Contact Data

There appear to be only a small number of tasks that a user can perform with the imported contact information. In particular:

- The user can send invitations to the imported contacts to become friends.
- The user can remove the imported data.

How the imported contact data is used by Facebook to make "People You May Know" suggestions was considered. Detailed technical documentation for the "People You May Know" functionality has been provided by Facebook and reviewed as part of this process. The imported contact may consist of both existing Facebook users and non-users. These were considered separately as follows:

- In the case of imported contact details of existing Facebook users, it was noted that:
    - Existing Facebook users in imported contacts may be used as the basis of "People You May Know" suggestions.
    - Other Facebook users who have imported the user as a contact may also, in some circumstances, be used as the basis of "People You May Know" suggestions.
- In the case of imported contact details of non-Facebook users, it was noted that:
    - As described in more detail in Section 3.5, the fact that two Facebook users have only a non-Facebook user imported contact in common does not appear to cause the two users to be suggested to each other as "People You May Know". This is consistent with the documentation provided by Facebook detailing the operation of the "People You May Know" functionality.
    - If multiple Facebook users have imported the same non-user email address, invitations sent to the non-user may contain as suggestions users that have imported the non-user's email address. Users who have already sent invitations to the non-user do not appear to be suggested in subsequent invitations.

While not being conclusive, it appears based on the above results, that the functionality by which Facebook users are suggested to each other as possible friends (referred to above as "People You May Know") and the functionality by which users are suggested to non-users in invitations operate on separate principles. Therefore, it seems likely that these two pieces of functionality are separate.

[1] This implies that if a single Facebook user imports a particular contact email address, and that user subsequently removes their imported contacts, then the reference to the imported contact will be removed from this structure.

2.4 Non-User Opt Out

When the non-user chooses to opt out of receiving subsequent invitations from Facebook, a hash of their email address is created and stored. A hash is a one-way function that generates a unique value representing a particular email address[2]. The key feature of these functions is that it is easy to calculate the hash value for the email address but effectively impossible to calculate the email address given the hash value. In this case, an MD5[3] hash of the non-user's email address is stored.

Certain scenarios can arise where other Facebook users perform an activity that would cause the non-user to receive email invitations. An example would be if a second user attempts to invite the non-user to join Facebook. The fact that non-user's email address matches an MD5 hash in the list of opted out email hash values will prevent the email from being sent.

Facebook were requested to provide a list of all the possible ways that a non-user of Facebook could receive an email from Facebook. The list provided was:

- A user invites a non-user to join Facebook
- A user sends a private message to a non-user
- A user creates an event and invites a non-user to the event

The impact of the user having opted out is that they will not receive any more invitations to join Facebook. The opted out non-user will also not receive invitations to events.

An opted out non-user will still receive private messages sent by users of Facebook. It was noted that private messages do not contain a link inviting the non-user to join Facebook.

2.5 Import Password

When importing contacts from an email account, the user will provide Facebook with the username and password of a supported email provider. Facebook will then use these credentials to connect to the email provider and import the contacts.

The code used to perform this functionality has been examined and it has been confirmed that the email provider password is stored in memory for the duration of the import task and then discarded.

---

[2] In some very remote scenarios, it is possible that multiple, specially crafted input values can lead to the same hash value. This is known as a collision. Regardless of the likelihood of such a collision, this is not considered relevant here and is only mentioned for completeness. The upshot of a hash collision in this case would be that a non-Facebook user who had not opted out of receiving emails would not receive invitations because the hash value of their email address matched the hash value of the email address of a non-Facebook user who had opted out. In particular, this would not lead to a situation where Facebook could recover non-user email addresses from stored hash values.

[3] http://tools.ietf.org/html/rfc1321

3. Synchronising

3.1 Background

Facebook provide a mobile platform for allowing users to interact with Facebook on their mobile devices. Applications are available for iPhone, Palm, Sony Ericsson, INQ, Blackberry, Nokia, Android, Windows Phone and Sidekick[4].

Testing has been performed on the iPhone version of the Facebook app (version 4.0.2). Only the contact synchronisation feature of the mobile application is under consideration as part of the present review.

The contact syncing functionality of the mobile application allows users of the application to synchronise the contacts in their address book with their Facebook friends.

3.2 Transmission of Contact Information

The information transmitted by the application while contact synchronisation is taking place has been captured and examined using the following process:

- A newly installed, fully patched Windows XP virtual machine with anti-virus software installed is used to create a new Facebook account. All browsing was carried out using the default configuration of Internet Explorer 8 (Version: 8.0.6001.18702).
- Two test contacts are transferred to the address book on an iPhone 3GS (iOS version 5.0.1). These contacts contained the following fields:
    - Name
    - Email
    - Phone
    - Address
    - Company name
    - Friend
    - Assistant
    - AIM
- Wireshark[5] is used to capture all traffic generated by the iPhone on the appropriate wireless network.
- The Facebook app is started and logged in to the new Facebook account.
- Contact syncing is enabled and the synchronisation is observed to take place.

The traffic generated in this way has been examined and the aspect of communication between the iPhone app and Facebook relating to the transmission of the contact information has been isolated. It was noted that the following data structure was transmitted to Facebook[6]:

[4] http://www.facebook.com/mobile/
[5] http://www.wireshark.org/
[6] The actual contact information has been removed in the excerpt provided. Names have been replaced with “<Contact 1 Name>, <Contact 2 Name>”, emails with “<Contact 1 Email>, <Contact 2 Email>” and phone numbers with “<Contact 1 Phone>, <Contact 2 Phone>”

```
[
        {
                “phones”:[“<Contact 1 Phone>”],
                “name”:”<Contact 1 Name>”,
                “emails”:[“<Contact 1 Email>”]
        },
        {
                “phones”:[“<Contact 2 Phone>”],
                “name”:”<Contact 2 Name>”,
                “emails”:[“<Contact 2 Email>”]
        }
]
```

It is therefore confirmed that only the contact name, email(s) and phone number(s) are transmitted to Facebook. None of the other contact information appears to be transferred.

It is not possible for the user to select a subset of contacts to synchronise with Facebook. The names, phone numbers and email addresses of all contacts in the phone's address book will be transferred.

3.3 Security of Transmitted Data

Facebook provides a range of optional security settings. One of which is known as secure browsing which will use a secure connection (https) to browse Facebook when possible[7]. This feature is not currently available for mobile browsing.

It has been confirmed that the contact information being synchronised is transmitted in plain text regardless of the state of the secure browsing setting.

3.4 Contact Synchronisation vs. Find Friends

The Facebook iPhone application has two closely related features, contact synchronisation and find friends. Both of these features are accessible by pressing the same button, in the top right hand corner of the "Friends" screen in the iPhone app.

Important distinctions in the behaviour of these features have been identified and are explained here. The following test was performed to illustrate the difference between the treatment of data using "Sync Contacts" and "Find Friends":

- A newly installed, fully patched Windows XP virtual machine with anti-virus software installed is used to create a new Facebook account. All browsing was carried out using the default configuration of Internet Explorer 8 (Version: 8.0.6001.18702).
- The Facebook iPhone app is installed on an iPhone
- The Facebook iPhone app is started and the new Facebook account is used to log in.
- The contact synchronisation feature is enabled.
- It was confirmed by logging in to Facebook that the contact information is not visible under "Manage Invites and Imported Contacts"

[7] http://www.facebook.com/help/?page=132501803490562

- “Find Friends” is clicked in the Facebook iPhone app.
- It was noted that the contact information is now visible under “Manage Invites and Imported Contacts”

It is notable that after contact synchronisation was enabled the contact information was not accessible from the user’s “Manage Invites and Imported Contacts” page[8]. The only way that was identified for interacting with synchronised contact information was via the Facebook iPhone app. It is important to emphasise that the “Find Friends” button in the iPhone app had not been pressed at this point.

Syncing can be disabled at any time through the iPhone app. The action of disabling synchronisation does not appear to delete any of the synchronised data.

There is a “Remove Data” button on the contact synchronisation screen in the iPhone app. There are two categories of data created by the synchronisation process that one would expect to be deleted;

- Data that was transferred from Facebook to the iPhone
- Data that was transferred from the iPhone to Facebook

The data that was transferred from Facebook to the iPhone constitutes the additional data that is added to the user’s address book, namely profile photos, birthdays and Facebook URLs. It has been confirmed that changes made to contacts by the contact synchronisation process are removed by clicking the Remove Data button.

The data that was transferred from the iPhone to Facebook, specifically the names, phone numbers and email addresses of all contacts in the phone address book, are not deleted from the Facebook servers when the “Remove Data” button is pressed.

It is possible to delete the synchronised contact information from Facebook’s servers, but it is not immediately apparent how to perform this task. Despite the fact that the synchronised contact information is not visible in the “Manage Invites and Imported Contacts” page, the “remove all your imported contacts” link will remove the synchronised contact data from Facebook’s servers. This has been verified as follows:

- Create a new test Facebook account and verify that no contacts are present in the “Manage Invites and Imported Contacts” link.
- Using Facebook’s internal tools, verify that no contact information is present in the user’s phone book.
- Install the Facebook iPhone app.
- Enable contact synchronisation in the iPhone app.
- Verify that contacts are not visible in the “Manage Invites and Imported Contacts” link.
- Using Facebook’s internal tools, verify that the contact information synchronised from the iPhone is present in the user’s phone book.
- Turn off contact synchronisation in the iPhone app.

[8] http://www.facebook.com/invite_history.php

- Via “Manage Invites and Imported Contacts”, use the “remove all your imported contacts” link to remove all imported contact information.
- Confirmed using Facebook’s internal tools that the contact information synchronised from the iPhone is no longer present in the user’s phone book.

It is not possible for a user without access to Facebook’s internal tools to verify that the synchronised contact information has been deleted. The fact that it is not apparent to the user how to manage their synchronised contact information is a shortcoming in the Facebook user interface.

As described below, the functionality of the “Find Friends” button appears equivalent to contact importation as described in Section 2. The fact that the synchronised contact data is not visible to the user without first performing a contact import (i.e. clicking “Find Friends”) will tend to blur the distinction between these two features both in reality and in terms of user expectations.

When the Find Friends button is clicked, the user’s address book information is presented in two categories in the iPhone app interface; firstly, any contacts that are existing Facebook users are displayed and the user can choose to send friend requests, and secondly, non-Facebook users are displayed and the user can choose to send invites to these contacts to join Facebook and become friends. Both sets of users are presented with an option to simultaneously send connection requests to all contacts being presented.

Only after the “Find Friends” button has been clicked is the contact information visible in the “Manage Invites and Imported Contacts” Facebook page. The Facebook user appears to be able to perform the same functions on the contacts that were available when contact importation was used to import the contacts. This functionality is described in Section 2.

The contact information can be removed via the “remove all your imported contacts” link on the “Manage Invites and Imported Contacts” page. This is exactly the same process that is followed when contacts are imported from any source with the proviso that removed contacts will be re-imported automatically unless you turn off syncing in the Facebook iPhone app.

Testing has been performed to confirm that after the contacts have been removed via the “remove all your imported contacts” page that they appear to be re-imported automatically from the iPhone app only after the “Find Friends” button is clicked.

As described in Section 2.2, a code review has been performed on the contact removal code to verify that the code actually deletes the imported contact information as expected.

3.5 Use of Non-Facebook Synchronised Contacts, Imported Contacts and Invites in “People You May Know” Calculations

A series of tests have been performed to attempt to understand how synchronised contacts, imported contacts and sent invitations relate to the generation of “People You May Know” suggestions.

If a Facebook user enables the contact synchronisation feature of the Facebook iPhone app, then if there are any existing Facebook users in the synchronised contacts, these will be suggested as

people you may know. The existing Facebook users are presented both as a separate list under "Find Friends" in the iPhone app and also may be presented in the "People You May Know" section of the Facebook web page.

The following scenario was considered:

- Two Facebook users that have no friends in common
- Both users install the Facebook iPhone app and enable contact synchronisation
- The two users have a non-Facebook user contact in common

The fact that the two Facebook users have a non-Facebook user contact in common does not appear to change the "People You May Know" suggestions for either user. In particular, the two users who have the non-Facebook user contact in common are not suggested to each other as "People You May Know". This appears to also be true if both of the users have sent invitations to the non-Facebook user (which the non-Facebook user has ignored).

Detailed technical documentation for the "People You May Know" functionality has been provided by Facebook and reviewed as part of this process. Non-user data is never reported as being used to generate "People You May Know" suggestions. The results of the testing described in this section are consistent with the documented functionality.

4. Data Security

This section on data security is divided into two sections; security of user communication with Facebook and Facebook corporate information security.

4.1 Security of User Accounts

Facebook provide a range of base security features by default on all accounts. The most obvious of these are the credentials used to log in. However, Facebook also monitor for suspicious activity on user accounts. Detection of suspicious activity will lead to additional authentication steps such as the user needing to fill out a CAPTCHA or by an SMS authorisation code sent to the user's mobile phone.

Facebook also provide a selection of opt-in security features, accessible under Account Settings->Security. They are:

- Secure Browsing
- Login Notifications
- Login Approvals
- Active Sessions
- One-time Passwords

Secure browsing enables the use of encrypted communication using HTTPS whenever possible. Secure browsing is not supported on the mobile platform. It has been confirmed that enabling secure browsing appears to causes all subsequent web browsing to be performed over HTTPS.

Login notifications involves notifying the user whenever their account is accessed from a computer or mobile device that has not been used before. Login approval involves entering a security code, which is sent to the user by SMS, each time the user's account is accessed from a computer or mobile device that has not been used before. It has been confirmed that enabling login notifications causes an SMS containing an authorisation code to be delivered whenever a login is attempted from a web browser from which the Facebook user has not logged in before. It has also been confirmed that it does not appear to be possible to log in without the authorisation code.

Active sessions allows a logged in user to see the locations from which their account is currently logged in and end activity from any particular session if that activity is unrecognised. It has been confirmed that ending activity in active sessions immediately causes the relevant user session to be logged out.

One-time passwords is a feature to allow users protect their account when they log in from a public computer. The user sends an SMS to a particular number and they will receive an eight character temporary password, valid for 20 minutes, which can be used to access their account.

The availability of the one-time passwords feature appears to depend on country and mobile operator.

Facebook maintains a security centre to provide a resource to educate users about staying safe online and maintaining the security of their account[9].

4.2 Corporate Information Security

A review has been performed of Facebook's corporate information security arrangements.

From the review it has been concluded that Facebook has appropriate information security controls in place, broadly consistent with the requirements of ISO 27001 and 27002.

The majority of the controls described by Facebook appear to be effective. If large-scale, frequent data breaches were taking place on Facebook's corporate networks, it is believed that this would be widely reported, particularly considering Facebook's global profile. Since this is not the case, the information security controls in Facebook appear to be preventing these types of incidents.

If there is a shortcoming in Facebook's information security arrangements it is their informality. Many policies and procedures that are in operation are not formally documented. The absence of these documented policies and procedures means that it is difficult to assess the audit trail data stored by Facebook within the context of their information security policy.

In the particular case of employee access to Facebook user data, Facebook retain a log of every access by every employee of every Facebook user account. This data is examined both automatically to identify patterns of suspicious behaviour and manually when specific cases require investigation. Facebook has demonstrated the functionality of their automated investigative tools. Facebook has also demonstrated that new abuse scenarios are added to the automated investigative tools as they are identified.

Particular attention was paid to the controls surrounding employee, contractor and vendor access to Facebook user data. The following items were noted:

- All employees, contractors and vendors are subject to the information security policy, and are required to familiarise themselves with the terms of the policy on a regular basis.
- Regular, company-wide security awareness training is carried out.
- Employees, contractors and vendors are required to sign a non-disclosure agreement before access to user data is granted.
- Contracts with third parties contain security and privacy requirements and periodic reviews of third party compliance with these requirements are carried out.
- A due diligence process exists that is used to assess if a third party has the capability to comply with the security and privacy requirements.
- Pre-employment screening is performed on all employees.
- An identity management system has been deployed to provision accounts, remove accounts and manage access rights.
- All users are assigned a unique user name and password. Password policy requirements are enforced on all systems.
- Credentials required to access production systems automatically expire on a regular basis requiring a manual process to re-enable access.

[9] https://www.facebook.com/security

- A manual process is required to grant an employee access to Facebook user data. The process requires approval by the data or system owner.
- Currently access rights are tool based, meaning that an employee with access to a particular tool can access any user data accessible through that tool. A new, software token-based access management system is under development to enable more fine-grained access control to user information.
- A valid certificate of PCI DSS[10] compliance pertaining to the storage of customer financial data has been presented.

[10] Payment Card Industry Data Security Standard. See https://www.pcisecuritystandards.org/security_standards/

5. Application Development

5.1 Background

Facebook provide an application platform to allow third party developers to build applications that integrate with the Facebook platform[11]. Facebook also provide development platforms for integration with other websites (e.g. social plugins which are discussed in Section 6) and integration with mobile applications.

The applications that form the basis of the testing described below are applications that integrate directly with the Facebook platform. These applications conform to the following basic architecture:

- Facebook applications are loaded into a canvas page that is populated by the third party application. An example of the URL of a canvas page would be https://apps.facebook.com/SimpleTestApplication/. This is the URL through which the user interacts with the application.
- The third party developer provides a URL, known as the canvas URL. Facebook submits requests to the canvas URL in order to retrieve content for presentation to the user on the canvas page.
- The content retrieved from the canvas URL is loaded within an iframe on the canvas page.

Facebook submits information about the user of the third party application to the canvas URL in the form of a HTTP POST with a single parameter called signed_request. This parameter is a base64[12] encoded JSON[13] object that must be decoded before processing.

The test applications described here were developed in PHP5[14]. For the simpler test applications, and where the technique was appropriate, direct querying of the Facebook APIs was performed. Facebook also provide a PHP SDK (Software Development Kit)[15] that was used in a number of the tests to simplify the development task. The code of the Facebook PHP SDK was reviewed and the relevant aspects were confirmed to operate as reported.

5.2 Application Access Control

Application access to user account information is controlled by permissions. The application must request permission to gain access to various types of information or perform actions on the user's account[16].

The minimum amount of access that a user can provide an application with will allow that application to access to their basic information. The basic information is:

---

[11] https://developers.facebook.com/
[12] http://tools.ietf.org/html/rfc4648
[13] http://tools.ietf.org/html/rfc4627
[14] http://www.php.net/
[15] http://developers.facebook.com/docs/reference/php/
[16] Facebook report that the application authorisation process is protected by Cross-Site Request Forgery (CSRF) checks to ensure that flaws in applications cannot cause a user to authorise an application without their knowledge.

- User ID
- Name
- Profile picture
- Gender
- Age range
- Locale
- Networks
- List of friends
- Any other information the user has made public

Access to other information about the user or their friends requires that the application request extra permissions from the user.

A complete list of permissions is available on Facebook's developer portal[17].

After having authorised an application, the user can revoke the authorised permissions through their account settings[18]. Certain types of permissions are required and can only be revoked by de-authorising the application entirely. Other permissions can be revoked individually. The list of permissions on the Facebook developer portal provides information on which permissions fall into each of these categories. The access previously granted to any application can be removed entirely by removing the application through this interface. It has been confirmed that an application that has been removed through this interface is no longer able to access any user information.

5.3 Before Authorisation

Before a user authorises an application to access any of their information, the application can access the country, locale and age range of the user.

This has been verified by creating a test application that does not request any authorisation from the user and displays whatever information Facebook provides when the user visits the application's canvas page.

When the signed_request value is decoded, the user information contained therein was:

- country = "ie"
- locale = "en_US"
- age = array{(min = 21)}

The age value is an array containing a min value. A max value may also be present, although one was not present in this case. These values represent the age range of the user.

[17] http://developers.facebook.com/docs/reference/api/permissions/

[18] Accessible via Account Settings->Apps and also through Privacy Settings->Apps and Websites->Apps you use->Edit Settings.

These parameters are provided so that an application developer can ensure that the content delivered by the application is appropriate for the age and country of the user, and is also localised appropriately.

The content of the HTTP request headers received by the canvas URL from Facebook were examined to ensure that the HTTP headers do not contain any user identifying information. In particular, it has been verified that the HTTP referrer header does not contain the user ID of the browsing user.

5.4 Application Authorisation

When the user has authorised an application to access their account according to a particular set of permissions, the application is provided with an authorisation token. This token is then provided to Facebook along with subsequent requests for information. All testing has been performed based on the server side authentication flow[19].

It has been confirmed that only basic information, as described above, is accessible when no specific permissions are requested by the application.

Exhaustive testing of all permissions has not been performed, however several sample permissions were selected and it has been confirmed that these permissions are required as expected to perform the actions governed by those permissions. In particular:

- If the application has not been granted the “user_photos” permission, a request to view the content of a user album that is shared only with friends, fails.
- If the application has been granted the “user_photos” permission, a request to view the content of a user album that is shared only with friends, succeeds.
- If the application has not been granted the “publish_stream” permission, an attempt to post a message to the user’s wall fails.
- If the application has been granted the “publish_stream” permission, an attempt to post a message to the user’s wall succeeds.

5.5 Access to User Friend Information

When a user authorises an application, that application can request access to the same information about that user’s friends as the user has access to. This access is not granted by default and must be specifically requested by the application. For example:

- User A starts using an application.
- User A authorises the application to access their friend’s photos.
- User B is a friend of user A.
- User B has some photos that are only shared with friends.
- User B has not indicated via privacy settings that their photos should not be shared with applications that their friends use.
- The application will have access to User B’s photos.

[19] http://developers.facebook.com/docs/authentication/

Unless the friend has opted out as described below, the same basic information listed in Section 5.2 is also available to the application about each of the user's friends.

Users can control what information the applications that their friends are using can see about them. This configuration is carried out in Privacy Settings->Apps and Websites in the section titled "How people bring your info to apps they use". The user can unselect any of the aspects of their profile that they do not want shared (except basic information). It has been confirmed that if, in the above example, User B has unchecked "My Photos" in this privacy configuration, an application installed by User A can no longer see User B's photos. This is true even if User A has authorised the application to view their friend's photos.

Note that User A will still be able to interactively view User B's photos by browsing to User B's profile, but applications installed by User A will not.

If a user does not want any information shared with applications that their friends install, including basic information, they must disable the application platform. This is achieved by selecting "Turn off all platform apps" in Privacy Settings->Apps and Websites. It has been confirmed that if the user decides to do this, applications that their friends install will have no visibility of them, including basic information. However, when the application platform is disabled the user will not be able to use any applications themselves.

5.6 Duration of Validity of Token

By default, the tokens generated by authorisation requests are valid for a period of 2 hours.

However, if the authorisation token request includes a request for the "offline_access" permission, the token returned upon approval will be valid for a longer period of time. The purpose of this permission is to allow applications to perform actions on the user's behalf at any time.

Tokens granted with the "offline_access" permission do not expire after any period of time. Rather, the token is invalidated based on the occurrence of certain events such as the user changing their password or suspicious activity on the user's account.

5.7 Transferability of Authorisation Token

It appears to be possible to generate authorisation tokens that authorise the bearer of the token to perform the action governed by the permissions requested. In the case described here, the validity of the token is not dependent on the source of the request. This has been tested as follows:

- Create a test Facebook user account.
- Create a test photo album, shared only with friends.
- Create a test application (TestApp1).
- TestApp1 requests only access to basic information.
- The user authorises TestApp1.
- Confirm that TestApp1 cannot view the non-public user photos.
- Create a second test application (TestApp2).
- TestApp2 requests additional permissions (user_photos permission in this case).

- The user authorises TestApp2.
- Confirm that TestApp2 can view the non-public user photos.
- Transfer authorisation token received by TestApp2 into application TestApp1.
- Confirm that TestApp1 can now view the non-public user photos.

The token also appears to be valid when used outside the context of the Facebook app platform. This has been tested in the following way:

- Create a test Facebook user account.
- Create a test application (TestApp1).
- TestApp1 requests additional permissions (publish_stream permission).
- The user authorises TestApp1.
- Confirm that TestApp1 can post messages on the user's wall.
- Create a static HTML file stored locally on a PC containing a form that performs a HTTP POST to submit the authorisation token received by TestApp1, a link and a message to be posted on the user's wall via the Graph API to the feed of the user ID who authorised TestApp1.
- Confirmed that submitting the form in the static HTML file causes a message to be posted on the user's wall.

In all cases, the authorisation token used for testing was the oauth_token value returned by the OAuth dialog[20].

Only a matching App ID and Canvas URL were required to generate valid authorisation tokens. The application secret did not appear to be required.

Facebook confirm that the architecture of the application authorisation/permissions system means that possession of the token authorises the bearer to access the information or perform the activities authorised by that token. The alternative to such as system is to require the application to generate some form of digital signature, based on the application secret, for each submitted request. Two reasons were provided for why Facebook moved away from this architecture to the current bearer token model:

- The greater complexity of the code required to sign requests meant that certain application developers were unable or unwilling to develop applications that used such a system. Facebook report that there is a strong correlation between the ease of use of an API and the uptake by the development community.
- Only one use case has been considered above, based on the server side authentication flow. Other use cases such as the use of Facebook APIs in Adobe Flash applications or standalone executable files have not been considered. In these cases, requiring the application developer to sign requests to Facebook often leads to the application secret being coded into the Adobe Flash or standalone application. It is then possible for a malicious individual to reverse engineer the application and retrieve the application secret. This security outcome is considered worse than the risk presented by the bearer token model.

[20] http://developers.facebook.com/docs/authentication/

It has not been possible to generate authorisation tokens for applications outside the control of the test developer account. The App ID and Canvas URL for any application can be retrieved by authorising access to the application to a Facebook account. However, it has been confirmed that control of the Canvas URL is also required since the Facebook response containing the authorisation token will be delivered to the Canvas URL.

5.8 Reliance on Developer Adherence to Best Practice/Policy

Several scenarios have been identified that require developer adherence to best practice or stated policy in order to ensure security of user data.

5.8.1 Use of Secure Site to Host Application

When a user authorises an application, the authorisation token is submitted to the canvas URL provided by the developer when the application was set up.

As of 1st October 2011 Facebook require the authors of applications to provide both a canvas URL and a secure canvas URL (i.e. a HTTPS URL). At the date that the testing was performed, around 3rd December 2011, it was not necessary to provide a secure canvas URL. A test application was configured with only a canvas URL and authorisation tokens were successfully delivered to this (unencrypted) URL.

This appears to introduce the risk that unless the application developer provides a secure canvas URL, authorisation tokens could be intercepted in transit to the application.

5.8.2 Cross-Site Request Forgery (CSRF)

Cross-site request forgery is an attack where a legitimate user visits a malicious website which causes an action to be performed on the user's account without the user's knowledge.

The OAuth 2.0 standard[21] upon which the Facebook application authorisation framework is based allows the transmission of an opaque state parameter that is returned to the caller along with the authorisation token. An application developer can use this feature to, amongst other things, ensure that the authorisation token is received in response to a known authorisation request. This technique reduces the risk from CSRF attacks.

Facebook strongly recommend in their application authentication documentation that any applications implementing Facebook user login implement CSRF protection using this mechanism.

5.8.3 Storage of Access Tokens

With reference to Section 5.7, there appears to be a significant risk associated with theft of authorisation tokens from an application developer. This risk is particularly acute in the case where authorisation tokens with long periods of validity have been generated (i.e. where the "offline_access" permission has been granted).

The testing in Section 5.7 would suggest that the compromise in a third party developer of valid pairs of user ID and authorisation token would allow any application to gain equivalent access to

[21] http://tools.ietf.org/html/draft-ietf-oauth-v2-22

user accounts as the access granted to the original application for the period of validity of the authorisation token.

It is, therefore, necessary that application developers take appropriate steps to ensure the security of the authorisation tokens provided by Facebook.

Facebook has the ability to suspend an application's access to the application platform, as well as systems to detect inappropriate actions and automatically disable suspicious applications. Referring again to Section 5.7, the suspicious actions will appear to Facebook as if they are being performed by the legitimate application from which the authorisation tokens were stolen since this is the application to which the authorisation token was granted by Facebook.

While disabling the legitimate application may indeed prevent the use of stolen authorisation tokens, it may also interrupt service for many legitimate users.

5.8.4 Increasing Access

Based on the permissions structure described in this section, an application that does not have access to a user's private information cannot increase the access to that user's information. In other words, a technical infrastructure prohibits unauthorised access to user data by applications.

However, an application with access to some of a user's private data could, hypothetically, increase the access that one user has to another user's data beyond what that user's privacy settings would allow. For example:

- User A only shares photos with friends.
- User B is not a friend of User A.
- User A authorises access to their photos to an application.
- User B authorises access to their photos to an application.
- The application has access to User A's photos and were the application to present User A's photo's to User B for any reason, User B would have gained increased access to User A's information.

Applications that increase access to information in this way are prohibited by policy[22].

It appears, *prima facia*, extremely challenging to implement a technical solution to ensure that applications do not perform this type of action.

[22] http://developers.facebook.com/policy/

6. Social Plugins

6.1 Background

Social plugins are a feature provided by Facebook to website owners, allowing the owners of websites to provide a customised browsing experience for Facebook users. The social plugin allows users to see relevant information such as which of their friends have "liked" the content of the website.

When a logged-in Facebook user visits a website that has a Facebook social plugin, the user will be presented with personalised content based on what their friends have liked, commented or recommended upon on the site.

6.1.1 Cookies

HTTP is a stateless protocol. This means that a web server does not retain information or status about the relationship between multiple requests. This represents a challenge in some use cases.

For example, if a website has a publicly accessible area and a private area to which authorised users must log in, the application running on the web server needs to track the fact that a particular user has logged in in order to serve information from the private area to only authorised users. A common technique used to achieve this goal is to provide a piece of information, known as a cookie, to the user's web browser. This piece of information is then returned to the server in every subsequent HTTP request, allowing the web server to associate the current HTTP request with the logged in user.

6.2 Social Plugin Structure

Social plugin content is loaded in an inline frame, or iframe. An iframe allows a separate HTML document to be loaded while a page is being loaded. In this case, the social plugin content is loaded separately from the content of the surrounding website.

It has been confirmed that the content of the social plugin component of a web page is delivered directly to the web browser from Facebook, separately from the surrounding content from the website as follows:

- 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).
- The website http://www.imdb.com/ was visited. This site contains a social plugin.
- While the site was being loaded, all traffic generated by the virtual machine was captured by an instance of Wireshark running in the host operating system.
- By right clicking on the area of the web site containing the social plugin and selecting "View Source" it is possible to view the HTML content of the social plugin.
- It can be confirmed that this content was delivered by Facebook by examining the captured Wireshark data in the following way:
    - Examine the DNS requests using the display filter "dns".
    - Note the DNS request for www.facebook.com and note the IP address to which the domain resolves (in this case 69.171.242.14).

- Examine the HTTP requests/responses to/from this IP address using the display filter “http and ip.addr==69.171.242.14”.
- Examine the content of the HTTP response and note that the HTML is the same as the content of the social plugin iframe.

This confirms that the content of the social plugin iframe was delivered directly to the web browser from Facebook. Web browsers do not allow cross-frame communication or access to data served from different domains so the website on which the social plugin is hosted does not have visibility of the content of the social plugin delivered to the web browser.

6.3 Non-Facebook Users and Cookies

Several experiments have been performed to determine what, if any, cookies are set and/or transmitted when a non-Facebook user visits websites with social plugins.

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).

Two techniques have been used to determine the cookies transferred to and from Facebook when browsing websites with social plugins:

- Examining the content of C:\Documents and Settings\<User>\Cookies, where Internet Explorer stores cookie files.
- Examining the network traffic generated by the virtual machine, extracting and examining the HTTP traffic.

Cookies, or equivalent persistent browsing data, can sometimes be found in other locations on a computer. For example, Adobe Flash Local Shared Objects (LSOs) can be used for purposes equivalent to a cookie. These possibilities were identified but by examining the HTTP traffic in this case, it was possible to confirm that only browser cookies need to be considered.

Two scenarios have been found to produce different results;

- A non-Facebook user who has never visited the www.facebook.com web page.
- A non-Facebook user who has visited the www.facebook.com web page.

To examine the case of a non-Facebook user who has never visited the www.facebook.com page, a period of browsing was carried out in a test virtual machine, as described above. Care was taken not to visit the Facebook page. Some of the websites visited had social plugins and some did not.

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. In total, 39 HTTP requests to www.facebook.com were generated over the course of approximately ten minutes of browsing.

The HTTP request/response pairs were individually examined. No Set-Cookie headers were received in any HTTP response from Facebook[23]. Cookies can also be created in other ways, for example by JavaScript, but examination of the HTTP requests confirms that no cookies, howsoever created, are transmitted to Facebook in any HTTP request.

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). No Facebook cookies were identified as being present in the Cookies folder.

Therefore, in this case, where the non-Facebook user has never visited www.facebook.com, no cookies are sent either to or by Facebook when a user visits websites containing social plugins.

To consider the case of a non-Facebook user who visits www.facebook.com, a new virtual machine configured as described above was used. All traffic generated was, once again, captured by an instance of Wireshark running in the host operating system.

When a non-Facebook user visits the www.facebook.com home page, three Set-Cookie headers are present in the response from Facebook. These headers lead to the setting of three cookies:

- datr
  - This cookie is set with an expiry time of 2 years.
  - The path of the cookie was "/" and the domain was ".facebook.com". This means that the cookie will be returned with all requests for domains ending ".facebook.com".
  - The value for the cookie provided was "JAzNTotT3EBIP2XVJYlVYxHw".
- reg_fb_gate
  - This cookie does not have an expiry time. This means that the cookie is known as a session cookie and will exist until the web browser exits.
  - The path of the cookie was "/" and the domain was ".facebook.com". This means that the cookie will be returned with all requests for domains ending ".facebook.com".
  - The value for the cookie provided was "http%3A%2F%2Fwww.facebook.com%2F"
- reg_fb_ref
  - This cookie does not have an expiry time. This means that the cookie is a session cookie and will exist until the web browser exits.
  - The path of the cookie was "/" and the domain was ".facebook.com". This means that the cookie will be returned with all requests for domains ending ".facebook.com".
  - The value for the cookie provided was "http%3A%2F%2Fwww.facebook.com%2F"

In total 12 HTTP request/response pairs to www.facebook.com were generated over approximately five minutes of browsing to non-Facebook websites, some of which had social plugins and some did not. These HTTP request/response pairs were individually examined.

[23] The Set-Cookie HTTP response header is used to send cookies from the server to the user agent. See http://tools.ietf.org/html/rfc6265#section-4.1.

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

The newly deployed 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.

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.

A test was performed to verify the deletion of unexpected and out of context cookies as follows:

- The testing was performed on a newly installed, fully patched Windows XP virtual machine with anti-virus software installed.
- Mozilla Firefox[34] version 8.0.1 was installed.
- The Tamper Data Firefox plugin[35] version 11.0.1 was installed.
- The website http://www.facebook.com/ was visited. The response was viewed in Tamper Data and it was noted that the response sets the datr, reg_fb_gate and reg_fb_reg cookies as described above.
- The website http://www.facebook.com/ was reloaded and the HTTP request was intercepted using Tamper Data.
- Two additional cookies were added to the HTTP request with values provided as follows:
    - c_user=something
        - A Facebook cookie provided in the incorrect context. This cookie is only appropriate in the context of a logged in user.
    - _b lib=blob
        - A non-Facebook cookie
- The HTTP response was examined and it was noted that both the c_user and blib cookies were unset by the server.

Facebook have highlighted certain scenarios in which the cookie management framework will not clear cookies. These are:

- Cookies with invalid names will not be cleared. Facebook does not set any cookies with invalid names.
- Cookies that Facebook believes will not be cleared upon request (e.g. data in the form of cookies inserted into the cookie header by mobile carrier WAP gateways).
- It is possible for a user to manually craft a cookie in their browser that will be sent to Facebook which Facebook is unable to clear because the parameters used to set the cookie (e.g. cookie path) are not known.

Finally, due to the asynchronous nature of Facebook's architecture, it is possible that a response is sent to the user before the cookie checks have been completed or before the login state of the user is known. In these cases, the cookie management will occur on the next appropriate request.

[34] http://www.mozilla.org/
[35] https://addons.mozilla.org/en-US/firefox/addon/tamper-data/

6.7 Non-Cookie Information

Aside from the cookie information described in the previous section, the other information that can be transmitted along with requests for social plugins are:

- The HTTP headers sent by the web browser to the web server[36]. Typically these will include:
  - The Accept header: content formats that the web browser can accept.
  - The Accept-Language header: content languages that the web browser can accept.
  - The User-Agent header: typically contains the type of browser software and the operating system.
  - The Accept-Encoding header: whether the web browser can accept compressed responses, and in what formats.
  - The Host header: The hostname for which the HTTP request is being made.
  - The Connection header: allows the sender to specify options that are desired for the particular connection. For example whether to keep open or close the connection after this request has been processed.
  - The Cookie header: contains the cookie values.
- Time and date of request
  - The time and date that the Facebook server received the request.
- Browser IP addresses
  - Performing a HTTP request involves setting up a connection between the PC on which the web browser is running and the Facebook server that will process the request. Establishing such a connection requires that the server must know the IP address being used by the client[37].

6.8 Logging

Facebook were asked to provide a list of all queries run against social plugin logs over a period of a month. These queries were analysed to assess the nature of the queries performed and in particular to determine whether any queries for the activity of individual users were performed.

A spreadsheet with almost 3,000 queries was reviewed and the following points were noted:

- Only a single query was identified containing individual object IDs. Each of the individual IDs in this query were investigated and all turned out to be IDs of Facebook pages.
- For efficiency purposes, views[38] of the raw data set are constructed in order for employees to more efficiently query certain portions of the social plugin logs. All queries identified in the spreadsheet that queried such views resulted in aggregate data being returned.

---

[36] What follows is intended only as a summary of the typical information provided in the typical HTTP headers. Full details on each of the provided headers, and a list of all possible headers can be found in the HTTP standard at http://www.w3.org/Protocols/rfc2616/rfc2616-sec14.html

[37] Certain scenarios exist, notably the use of NAT (Network Address Translation) or the use of a web proxy, where the browser is not making a direct TCP/IP connection to Facebook. In these cases the IP address received by Facebook will not necessarily be the same IP address as that of the browser's PC.

[38] http://en.wikipedia.org/wiki/View_%28database%29

Although these results are not conclusive, there is no evidence in the information presented that individual user or non-user browsing activity is being extracted from social plugin logs for analysis.

6.9 Use of Social Plugin Activity to Target Advertising

Experiments have been performed to attempt to establish whether interacting with websites containing social plugins would influence the advertising that the user is presented with.

The first test performed was as follows:

- A target interest was identified, in this case “parenting/childcare”
- A Facebook user account with no previous browsing activity relating to this target term was used for testing.
- Browsing to various parenting/childcare websites with social plugins was carried out. This activity did not appear to have any discernable impact on the presented advertisements, as measured by noting the advertisements presented in the “Sponsored” column on the right hand side of the Facebook profile page.
- Different advertisements are presented over a period of time, but nothing that appears to correlate with the browsing activity.
- The test was repeated using a different test account and a different target term and the result was the same. The target terms used for the repeat test were “motorcycles”. No notable correlation was observed between the target terms and the advertisements presented to the user.

A second test was performed as follows:

- A second period of browsing was performed to parenting/childcare websites.
- This time, any available “Like”, “Share” or “Recommend” buttons were pressed.
- It was noted that a list of parenting/childcare websites now appears in the “Activities and interests” section of the Facebook profile.
- This activity also did not appear to have any discernable impact on the presented advertisements, measured as described above.
- Different advertisements are presented over a period of time, but nothing that appears to correlate with the browsing activity.
- This test was also repeated using a different test account and a different target term (“motorcycles”) and the result was the same. No notable correlation was observed between the target terms and the advertisements presented to the user.

A third test was performed as follows:

- Specific target terms were identified, in this case “Harley Davidson”.
- A Facebook user account with no previous browsing activity relating to this target term was used for testing.
- The Facebook pages of the Harley Davidson company and a specific Harley Davidson dealership were liked.
- The user’s location was configured as being in the geographic vicinity of the selected Harley Davidson dealership.

- Within 30 minutes, it was noted that advertisements relating to Harley Davidsons and motorcycles were targeted at the user.
- The test was repeated using a different test account, a different target term ("Cisco") and a location of Oakland, California. The result was the same. Within a short period of time advertisements relating to Cisco certifications were targeted at the user.

A fourth test was performed as follows:

- A target interest was identified, in this case "Cisco"
- A Facebook user account with no previous browsing activity relating to this target term was used for testing.
- The user's location was set to Oakland, California.
- Browsing to various Cisco related websites with social plugins was carried out. This activity did not appear to have any discernable impact on the presented advertisements, measured as described above.
- Different advertisements are presented over a period of time, but nothing that appears to correlate with the browsing activity.

A fifth test was performed as follows:

- A target interest was identified, in this case "Cisco"
- A Facebook user account with no previous browsing activity relating to this target term was used for testing.
- The user's location was set to Oakland, California.
- The website www.cisco.com was visited, the social plugin was located and the "Like" button was clicked.
- Within a short period of time advertisements relating to Cisco certifications were targeted at the user.

From the above tests, the following conclusions were drawn:

- The act of browsing to websites containing social plugins does not appear to have any influence on the advertising targeted at the user.
- Pressing the "Like" button either on a Facebook page or on a page with a social plugin may influence the advertising targeted at the user.
- The advertising targeting appears to be focussed on particular Facebook pages and/or very specific keywords.
- Behavioural profiling was not evident in the results described above. Browsing to a category of websites or interests (e.g. "parenting/childcare" or "motorcycles") did not appear to have any influence on the advertising targeted at the user. It is possible that other advertisements may use broader categories or keywords for targeting, but such ads were not identified in this testing.

Care is required when attempting to reproduce these results due to the very specific nature of the targeted advertising. As well as keywords being targeted, factors such as the geographic location, age and gender of the user may be considered by the ad targeting.

7. Akamai Cache

7.1 Background

In order to facilitate faster loading of the Facebook page, static content such as images and JavaScript files are cached using the Akamai caching service. Akamai maintain a globally distributed network of cache servers that will store copies of content on servers geographically closer to the users of that content than the source servers.

In this particular case, Facebook's data centres are located in the United States and users in locations far from the source servers benefit in terms of user experience when the static content is loaded from Akamai servers that are geographically closer to them.

7.2 Facebook Akamai URL Structure

A representative example of the filenames given by Facebook to image files is:

https://fbcdn-sphotos-a.akamaihd.net/hphotos-ak-ash4/387755_115906941856985_100003130400274_87779_1581190684_n.jpg

The photo file names produced by Facebook consist of five numeric components, as follows:

1. The volume ID: The first number is an identifier for the physical Facebook server where the image is located.
2. The Facebook object ID: The second number is a unique identifier for the photo.
3. The User ID: The third number is the user ID of the user who uploaded the photo.
4. The Photo ID: The fourth number is a legacy photo ID. For newer photos, this value is ignored. Specifically, if a Facebook object ID is provided, the Photo ID is ignored.
5. The fifth number is a pseudo-random number between zero and $(2^{31} - 2)$

The source code used to generate the photo file names has been examined and has been confirmed to work as described above.

7.3 Randomness of the Random Number

During the period of time that the testing described here was being performed, Facebook changed the technique being used to generate the random number component of the photo file names. Both the old and the new techniques are described here for completeness.

7.3.1 Older Technique Based on PHP mt_rand()

The older technique for generating the random number component of the photo file name was based on the PHP mt_rand() function[39]. This function generates pseudo-random numbers using the Mersenne Twister[40] algorithm.

The Mersenne Twister algorithm is not cryptographically secure, since from a sufficiently long subsequence of the outputs, one can predict the rest of the output[41], [42].

[39] http://php.net/manual/en/function.mt-rand.php
[40] http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html
[41] http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/efaq.html

In the case of Facebook photo file names, a pseudo-random number is generated for the photo file name upon upload of the image to Facebook. It is highly unlikely to ever have been possible to gather a sufficiently long string of consecutive pseudo-random numbers allowing prediction of subsequent pseudo-random numbers. This would have involved uploading the required number of images to Facebook such that a sufficiently long sequence of consecutive uploads was performed without being interrupted by anyone, anywhere else in the world uploading an image.

Even if this had been possible, being able to generate subsequent or previous pseudo-random numbers is of questionable value since it would not have been easily possible to associate a particular pseudo-random number value with an arbitrary photo uploaded by a user at some time in the past or future. In other words, it was challenging to associate a particular pseudo-random value with a corresponding volume ID, Facebook object ID and user ID.

In order to generate the photo file name for a particular photo, it seems that the following information would have been required:

1. The volume ID of the target photo
2. The Facebook object ID of the target photo
3. The user ID of the user who uploaded the target photo
4. A sequence of 624 consecutive pseudo-random numbers, generated by uploading a series of 624 photos to at least the same Facebook server[43] as the target user's photo (without being interrupted by anyone, anywhere else in the world uploading a photo). This string of 624 numbers would, preferably, have been generated close in time to the upload of the target photo. Were it possible to generate such a sequence, it would then have been possible to generate forwards and/or backwards through the Mersenne Twister sequence. Each value of the Mersenne Twister sequence would need to be tried with the other three pieces of information until the valid photo name was found.

A brute force attack was therefore hypothetically possible, but extremely unlikely.

7.3.2 Newer Technique Based on openssl_random_pseudo_bytes()

The newer technique for generation of the random number component of the photo file name is based on the openssl_random_pseudo_bytes() function which can generate cryptographically strong random number sequences[44].

---

[42] Several sources on the Internet indicate that a sequence of 624 consecutive integers is sufficient to predict all subsequent output. For example http://jazzy.id.au/default/2010/09/22/cracking_random_number_generators_part_3.html

[43] It may have been necessary, depending on implementation details not explored, that all 624 uploads would need to have been processed by the same web server process, and possibly by the same web server thread. There are typically multiple web server processes running on a web server and it is not possible for a user of the site to control which of the processes will be used to handle an incoming request. The use of load balancing will also mean that the user does not have control of which web server will handle any particular incoming request.

[44] http://php.net/manual/en/function.openssl-random-pseudo-bytes.php

The cryptographic strength of the output from this function has been tested. 200 million bytes of random data was generated and the randomness of the resulting output was assessed using the dieharder suite of cryptographic tests[45]. The dieharder suite has 27 tests that are marked as reliable for use when assessing randomness[46]. The 27 reliable tests were repeatedly run on the random data generated by openssl_random_pseudo_bytes().

In total 113 individual tests were run on the generated data, of which 97 (85%) passed. This provides a strong indication that the data is truly random.

7.3.3 Conclusion

In order to be able to generate photo file names to which an individual did not previously have access, the volume ID, Facebook object ID, user ID and random number component are all required.

Based on the above analysis, it appears that both the older and the newer techniques generate pseudo-random numbers of sufficient strength that it is not possible to guess the random component of an arbitrary photo file name.

The simplest way to have possession of the volume ID, Facebook object ID and user ID corresponding to a particular image is to have viewed the image in a browser. In this case, the whole file name is known so the random number will be available and a brute force attack is not required.

In the case where an attacker does not have access via Facebook to a target image; even if it were possible to guess the value of the pseudo-random number, this will not help to recover the volume ID or Facebook object ID of the image, regardless of whether the target user ID is known or not.

In conclusion it is believed that, until it is positively demonstrated to be flawed, the process used by Facebook to create photo file names is sufficiently robust to prevent generation of arbitrary, valid photo file names to which an attacker did not already have access.

7.4 Deletion of Facebook Photo

After a user has deleted a photo, Facebook no longer provide the Akamai cache URL to that photo. To confirm that this is the case, a technique for causing the Akamai cache to query Facebook on each request was used. Consider the representative image URL provided above:

---

[45] http://www.phy.duke.edu/~rgb/General/dieharder.php

[46] Version 3.31.1 of the dieharder suite was used in this testing. The tests run were Diehard Birthdays Test, Diehard OPERM5 Test, Diehard 32x32 Binary Rank Test, Diehard 6x8 Binary Rank Test, Diehard Bitstream Test, Diehard Count the 1s (stream) Test, Diehard Count the 1s Test (byte), Diehard Parking Lot Test, Diehard Minimum Distance (2d Circle) Test, Diehard 3d Sphere (Minimum Distance) Test, Diehard Squeeze Test, Diehard Runs Test, Diehard Craps Test, Marsaglia and Tsang GCD Test, STS Monobit Test, STS Runs Test, STS Serial Test (Generalized), RGB Bit Distribution Test, RGB Generalized Minimum Distance Test, RGB Permutations Test, RGB Lagged Sum Test, RGB Kolmogorov-Smirnov Test, Byte Distribution, DAB DCT, DAB Fill Tree Test, DAB Fill Tree 2 Test, DAB Monobit 2 Test. All tests were run with their default parameters.

https://fbcdn-sphotos-a.akamaihd.net/hphotos-ak-ash4/387755_115906941856985_100003130400274_87779_1581190684_n.jpg

An arbitrary query string was added to the end of this URL. For example:
https://fbcdn-sphotos-a.akamaihd.net/hphotos-ak-ash4/387755_115906941856985_100003130400274_87779_1581190684_n.jpg?a=b

Since this URL is different, the original, cached, content will not be returned and Akamai needs to pass the request through to Facebook. When the Facebook servers receive the request, the query string component of the request will be ignored and the original image will be returned.

Based on the above technique, the following process is used:

1. An image is uploaded to the Wall of a Facebook user.
2. The image is viewed and the Akamai URL is saved.
3. The query string "?a=b" is appended to the end of the Akamai URL and it is confirmed that the same image is returned by the Akamai cache.
4. The image is deleted from the Wall of the Facebook user.
5. The query string "?c=d" is appended to the end of the original Akamai URL and it is confirmed that no image is returned.

The original image URL will continue to return the deleted photo for a period of time. Facebook report that the Akamai cache retains content for 30 days, after which point it is removed from the cache. This has not been verified.

Therefore, it can be concluded that once the user has deleted the image, Facebook will not provide the Akamai URL at which the deleted image is cached to anyone viewing the user's profile.

In order for an attacker to retrieve from the Akamai cache a picture that a user has deleted from their Facebook profile, the attacker must therefore have prior knowledge of the photo URL.

In such cases, to retrieve the photo URL from Facebook, the attacker will most likely have viewed the image from the user's profile in their browser. Therefore, they may also have copies of the image cached locally on their PC and/or transparently cached, for example, by their Internet service provider.

8. Scraping

Scraping, also known as screen scraping, is the name given to an automated process of harvesting data from a website. In the case of Facebook, the concern surrounds the ability of such an automated process to gather a large volume of information about Facebook users through the scraping technique.

Facebook have provided details of the arrangements that they have made to prevent scraping.

It is believed that the current arrangements adequately mitigate the risk of large-scale harvesting of Facebook user data while allowing service to be effectively provided to legitimate users in a wide range of circumstances.

9. Account Creation Cancellation

A user begins the Facebook registration process by populating the registration form with their name, email address, gender and date of birth. This information is submitted to Facebook and then the user is presented with a CAPTCHA. It is only on the CAPTCHA screen that the new user is presented with their first opportunity to read the Terms of Use and the Privacy Policy.

If the user chooses to cancel their account registration for any reason at this point, the data entered into the first screen of the registration process must be removed.

The user may be sent reminder emails during a 30-day period, asking them if they want to return to complete the registration process. After 30 days, if the user has not completed the registration process, an automated process will delete the information provided.

The code of this automated process has been reviewed and confirmed to operate as specified, deleting all information stored when the user filled in the first page of the registration.

10. Account Deletion

10.1 Background

Facebook users can choose to either deactivate or delete their account[47].

If a user deactivates their account, this means that the user's profile information will not be available on Facebook, effective immediately. However, Facebook retains the user's information indefinitely in case the user chooses to reactivate their account at some point in the future.

Deletion, on the other hand, leads to the permanent removal of the user account from Facebook. The process followed when the user requests that their account is deleted is described in the next section.

10.2 Deletion Process

After a user submits a request to delete their account, their account enters a state of "pending deletion" for 14 days. During these 14 days it is possible to change your mind and cancel the deletion. This 14 day period is provided for various reasons, including allowing the user a "cooling off" period and also for the case where someone with unauthorised access to a user account issues a delete instruction.

If the user logs in to their account during the 14 day period where the account is pending deletion, they are presented with a message stating "Your account is scheduled for deletion. Are you sure you wish to permanently delete your account?". The user can then either confirm or cancel the deletion process.

Once the 14 day period has expired, an account deletion framework is activated which deletes account information. This account deletion framework is discussed in more detail below.

When an account is deleted all internally (within Facebook) and externally (on the Facebook website) visible information about the user is removed. In particular all user generated content directly associated with the user's account and virtually all other data associated with the user's account (directly or indirectly) is removed.

There are a small number of exceptions to this deletion process, which are discussed in the following section. None of this remaining data is visible on the Facebook website after account deletion.

10.3 Information Not Deleted by the Account Deletion Process

10.3.1 The Fact That The Account Has Been Deleted

After the account has been deleted, the fact that an account with a particular user ID has been deleted is recorded. Specifically the following information is retained:

- The user ID of the deleted account
- The status of the account is deleted

[47] http://www.facebook.com/help/search/?q=how+do+i+delete+my+account

- The time and date when the account was deleted

This is recorded for several reasons:

- To allow Facebook to be able to distinguish between the case of a user ID that has never existed and a user ID that used to exist but has been deleted.
- To enable Facebook to re-run the deletion process on the account in order to remove additional information that was missed when the account was originally deleted.

10.3.2 Log Data

Due to the nature and volume of log data stored by Facebook, deletion of all log records associated with individual user accounts is not possible.

At the time of writing, this log data is not deleted. However, a process is being implemented to permanently remove the relationship between log records and the user ID being deleted, effectively making the log records anonymous. The process works as follows:

- Ninety days after a particular log record is created, the user ID associated with the log record is removed and replaced with a random ID. The random ID is associated with the user's account in a userID-to-randomID mapping. Other data that may identify the user such as the datr cookie value, IP address, email addresses and text of any user generated content are also removed at this point.
- When a user account is deleted, the userID-to-randomID mapping is also deleted. Without this mapping, there is no way to know which userID the randomID was associated with and the log records are permanently anonymised.
- Therefore, after the user's account has been deleted, Facebook can no longer retrieve any logs that were associated with that user's account. Nor is it possible to determine the user ID of the deleted account from any log entry formerly associated with the account.

Samples of log data have been reviewed in their original form and the rewritten form after the user ID has been replaced with a random ID and other identifying information has been removed. It has been confirmed that the new log rewriting functionality operates as described above.

The log rewriting process described here has been under development for over a year but considering the volume of data involved the change is not instantaneous. The estimated time to completion of this log rewriting process is six months. Upon completion, as well as being used for all future account deletions, the log rewriting process will be applied retroactively to all accounts that have been deleted prior to deployment.

10.3.3 Shared Content

Shared objects, such as groups, pages and events, present a challenge for account deletion.

If the owner of a particular shared object chooses to delete their Facebook account, there is a question of whether the shared object that they created should also be deleted. Consider the case of a group; there are several scenarios that can be identified where it is less than ideal to automatically delete the group when the creator of the group deletes their account. For example;

- The owner of the group may no longer be actively involved in the group, but the group may consist of a vibrant community of other users who would not want the group deleted.
- It is possible that the owner of the group and the other members of the group have fallen out and the user's deletion of their account is an act of maliciousness directed towards the other members of the group.

Similar scenarios can be identified for pages and events.

At the present moment, most shared objects are deleted when the owner or original creator of the shared object chooses to delete their account.

However, some group content may remain. In particular, any posts that the user has made to groups will not be deleted. It is computationally easy to identify the user that created a particular posting in a group, but given only the user ID it is very difficult to identify all group posts made by that user.

This is because there is only a one-way relationship stored in Facebook's data relating the user ID to the group post. This relationship allows the user's profile picture to be looked up and displayed beside the content of their group post.

A solution is currently being implemented to convert this one-way relationship into a two-way relationship, allowing all of the user's group posts to be efficiently identified and removed when the account is being deleted.

10.3.4 Data that Can Not be Easily Found

Two examples of this type of data have been provided:

- There is some data which is not directly associated with the user being deleted and is only associated with another user with which the user being deleted was interacting. In these cases, an exhaustive search of the entire Facebook user space would be required to identify and delete this data. An example would be a comment on another user's wall post or a posting to a group, as described above.
- There is some data that was only linked to the user being deleted via an intermediate object that has been deleted. An example would be certain types of request that contain an attachment. The request may have been deleted but due to a bug in the deletion process, the attachment remains.

There is no easy way for Facebook to find and delete this data. Facebook are working on these technical issues and have undertaken to retroactively delete this information when solutions become available.

10.3.5 Messages

Message deletion is performed based on reference counting[48]. This means that when the last user who had a copy of the message deletes the message, that message will be permanently deleted. In the case of account deletion:

- If the user being deleted was the last user that had a copy of the message, that message will be deleted.
- If any other users have a copy of the message, the message will not be deleted.

It has not been possible to positively confirm that messages are deleted when the reference count drops to zero.

10.3.6 Other Non-identifying Data

It is possible that small amounts of non-identifying data have remained after users have deleted their accounts in the past. One example of this category of data would be information in a database table used to relate two pieces of data to each other[49]. Both of the pieces of information to which the relationship refer have been deleted, but the relationship entry was not deleted.

Facebook have an active process underway to identify and retroactively delete such data, as described in Section 10.4.

10.4 Account Deletion Framework

Facebook have, over time, identified a number of issues in their account deletion processes. In particular,

- There has been no way to verify that a user's account information has been deleted.
- If, for any technical reason, the deletion process failed or crashed, there was no way to retroactively seek out and delete information that was no longer associated with any active account.

Facebook has deployed a new user account deletion framework that is intended to address these issues. A due diligence process is underway to exhaustively identify all locations where user data is stored and to ensure that

- All new account deletion requests delete all user data from all possible locations.
- The new account deletion framework is applied to all previously processed account deletion requests that may not have adequately purged user data from all possible locations.

[48] Reference counting is a technique of storing the number of references to a resource and using this reference count to deallocate objects that are no longer required. http://en.wikipedia.org/wiki/Reference_count

[49] A particular example would be a many-to-many relationship table. See http://en.wikipedia.org/wiki/Junction_table.

Facebook confirm that at the present moment, any data that is not deleted when the user's account is being deleted is definitely made inaccessible and cannot be associated with the deleted user.

10.5 Deletion Verification Tests

Three tests to verify the status of a deleted account were performed.

Facebook were provided with the email address and full name of a user who had requested that their account be deleted more than 14 days prior to the test date. Facebook had no prior knowledge of either the email address or the full name. They were asked to provide any information that was available on their systems relating to this email address or full name. This test was performed under supervision and notes were made of the activities performed.

No details relating to either the email address or full name were found. The process used to search for the email address and full name were repeated with known Facebook user email addresses and full names to verify that if the test account existed, the searching performed by Facebook would have revealed the account information.

Facebook were provided with an IP address and asked to produce any information relating to browsing activity originating from that IP address. Facebook had no prior knowledge of the IP address.

Originally it was expected that the search world be performed over a 90 day period, however the Facebook log querying interface can in principle, but cannot in practice, query such a large date range. For illustrative purposes, querying Facebook's logs to identify the activity associated with a particular IP address in any given 24 hours period takes approximately one hour. The period of the search was therefore reduced to eight days.

No browsing activity was identified as being associated with the provided IP address.

Facebook acknowledge that this is an unexpected result for any IP address that is being used actively for browsing. No additional information is known about the browsing patterns associated with the IP address, however there are a number of possible explanations for this results;

- Firstly, if the user of the IP address had enabled a web browser privacy plugin that would block Facebook social plugins from being displayed on third party sites[50], no logs for websites with social plugins would be present.
- Another explanation would be if the user's web browser were configured to use a proxy. In this case, the browsing activity would not appear to Facebook to be originating from the user's PC IP address.

A further experiment was performed to determine whether after an account has been deleted, no information about that account remains.

[50] Many such privacy plugins exist, but one example is http://www.ghostery.com/.

The test involved comparing 62 types of information about an active account with the corresponding information for a deleted account[51]. For the active account, many of these types of information were present. For the deleted account, none of these types of information were present. Information about the fact that the account had been deleted remained, as described in Section 10.3.1, and it is possible that other data that is not efficiently accessible remains, as described in the remainder of Section 10.3.

[51] About me, address, alternate name, applications, check-ins, connections, credit cards, credits balance, currency, current city, date of birth, education, emails, events, family, favourite quotes, friend requests, friends, gender, groups, home town, last location, linked accounts, locale, logins, logouts, machines, mini-feed, name, first, last, name changes, networks, notes, notification settings, notifications, password, phone numbers, photos, physical tokens, pokes, political views, privacy blocks, privacy settings, profile blurb, real time activities, recent activities, registration date, relationship, religious views, removed friends, screen names, shares, status updates, subscribed to, subscribers, unified messages, vanity, videos, wall posts, website, work.

Appendix 2

Summary of Complaints

Summary of Europe v Facebook Complaints 1-22

Complaint 1 – Pokes

A poke is a type of short message sent from one Facebook user to another. Complainant stated that while Facebook allows for the removal of old pokes, they are not, in fact, being deleted. As part of an access request the complainant was provided with a copy of all pokes ever sent or received going back over a 2 year period.

Complaint 2 – Shadow Profiles

Complainant contended that Facebook is gathering information in relation to non Facebook users. This information primarily consists of email addresses, but may also include names, telephone numbers and addresses. Facebook typically collects this information when a user synchronises their phone or imports their email contact list to their Facebook account. Complainant contended that Facebook is using this data to create profiles of non-users. The complainant also stated that, on foot of an access request, Facebook did not provide any details in relation to Facebook users who may have uploaded his email or telephone details to their Facebook account.

Complaint 3 – Tagging

Friends on Facebook have the facility to 'tag' photos of another user (friend) and display them on their Facebook page and within the 'news feed' section. While the other user may subsequently remove the tag, the issue is that the user is unable to prevent friends tagging photos to their Facebook page in the first place.

Complaint 4 – Synchronising

This is related to Complaint 2. When a Facebook user synchronises their mobile phone or other device with Facebook, the complainant states that all personal data on the device are transferred to Facebook. This may result in invitations being issued by Facebook to those individuals whose data have been transferred. The individuals are not aware that their personal data has been disclosed in this way.

Complaint 5 – Deleted Posts

Facebook provides a facility whereby a user can delete old posts. The complainant stated that, while the act of deleting a post does remove the post from view, it is not, in fact, deleted. As part of an access request, the complainant was provided with posts which he had deleted. The deleted posts in question go back over a period of 3 years.

Complaint 6 – Posting on other Peoples Pages

When posting a comment on another person's Facebook page, the information posted is subject to the other person's privacy settings. The issue for the complainant is that the person posting the comment is unaware of the other user's privacy settings and, accordingly, will not know who can see the comment being posted – it could be restricted to friends only, but equally, could be viewed by everyone on the internet, including search engines.

Complaint 7 – Messages

It is possible for Facebook users to send instant messages to other users who are online. It is also possible to delete these messages if the user so chooses. However, the complainant contended that the act of hitting the delete button provided merely removes the message from view and does not, in fact, delete it.

Complaint 8 – Consent and Privacy Policy

There are a number of aspects to this complaint set out under the following headings:

Access: the complainant stated that Facebook's Privacy Policy is not easily accessible – the link 'privacy' provided at the bottom of the user's Facebook page is merely a link to a privacy guide, containing limited information. There is a link within this document to the actual Privacy Policy.

Role of FB-I and the User: the complainant stated that the user is not provided with any clear information on who is data controller (Facebook Irl. or Facebook Inc).

Extent of Privacy Information: the complainant was dissatisfied that, in order to get a grasp of Facebook's privacy policies, a user must deal with multiple documents and links, with many specific provisions difficult to locate.

Contradictions: the complainant highlighted contradictions he has identified within the Privacy Policy. He stated that the contradictions identified run to 6 pages and has provided some sample issues in the complaint in relation to the deletion of data.

Vague Provisions: the complainant has highlighted a number of provisions in the Privacy Policy which he considered to be vague and general in nature.

Unambiguous Consent: the complainant has highlighted a number of issues with the process of consenting to the Privacy Policy including the use of small text and lack of a check box to be ticked.

Freely Given Consent: this complaint is in relation to the monopoly Facebook has on business and personal users and that there should be a high bar in terms of privacy terms and conditions given the limited competition.

Specific Consent: the complainant contended that there is no specific consent being provided by users for the use of their personal data.

Informed Consent: the complainant considered that the purpose for which personal data is processed is not properly explained.

Consent obtained by deception or misinterpretation: this again relates to how Facebook uses personal data and the complainant has highlighted a number of examples where he considered Facebook to be providing false or misleading information.

Complaint 9 – Face Recognition

In relation to Facebook's use of tagged photos, the complainant contended that there has never been a specific consent provided to users availing of this feature, and particularly in relation to users who would have provided consents to Facebook prior to the introduction of the feature.

Complaint 10 – Access Requests
This complaint relates to an incomplete access request. The complainant stated that his access request resulted in only limited data being provided. He outlined the areas - 19 of them - under which he contended Facebook did not provide information.

Complaint 11 – Removal of Tags
A user is provided with the option 'remove tag' from a tagged photo on their Facebook page. However, the complainant contended that removing the tag is not deleting the tag data and that Facebook is not transparent in terms of informing users on the retention of this information following the use of the 'remove tag' option.

Complaint 12 – Data Security
The complainant set out a number of security concern in relation to the security of personal data – no encryption on private data (other than passwords and credit cards), not taking enough responsibility for data security in its privacy statements and a lack of control over data being provided to third party applications.

Complaint 13 – Applications
This relates to third party applications which sit on the Facebook Platform. The complainant outlined a number of issues including a lack of informed user consent when accessing a third party application, a lack of oversight by Facebook in terms of privacy compliance among third parties and the non-notification to users by Facebook in a case where a third party has no privacy policy.

Complaint 14 - Removed Friends
Facebook provides a facility to add friends and to 'unfriend' friends. The issue here is that when a user clicks on the 'unfriend' option, the friend information is not deleted, but is retained in the background – the complainant saw no justification for the retention and considered that Facebook is not transparent in terms of informing users on the retention of the information.

Complaint 15 – Excessive Processing
This covers a number of earlier complaints in terms of the non-deletion of information (pokes, tags, etc.) which a user may have removed from their page. The complainant considered the amount of data Facebook holds and processes to be excessive and a security risk.

Complaint 16 – Opt Out
The complainant argued that there is no specific consent when signing up to Facebook (see check box issue in Complaint 8) and that personal data is being collected prior to a new user being able to set their privacy settings. The complainant also contended that the security settings themselves are too liberal in nature and that the settings pages and links provided discourage the new user from applying certain security settings.

Complaint 17 – Like Button

The complainant stated that when a user visits a website which contains a 'social plug in' – the Like button – the following information is being recorded: date, time, URL, IP address, browser and operating system information. The complainant considered that the information is being collected unfairly and is excessive.

Complaint 18 – Obligations as Processor
The complainant considered that FB-I is a data processor while the Facebook user is the data controller. He contended that Facebook's operation as a processor is at variance with both Irish Data Protection legislation and Directive 95/46/EG.

Complaint 19 – Pictures Privacy Settings
Facebook allows a user to upload photographs to their Facebook page and apply security settings. The complainant stated that Facebook has outsourced the delivery of the picture content to a company (Akamai Technologies) and, by using the source code from the pictures page of Facebook.com and identifying certain URLs, that it is possible to view some photos that should be hidden from view.

Complaint 20 – Deleted Pictures
This complaint relates to the previous complaint (19). Facebook users are given the option to delete pictures they have uploaded to Facebook. Again, by using the source code from the pictures page of Facebook.com and identifying certain URLs, the complainant stated that it was possible to view a photograph for up to 48 hours after he had deleted it from Facebook.

Complaint 21 – Groups
Facebook allows users to add friends to groups which are found on the user's Facebook page and within 'news feeds'. The issue raised in the complaint is that a user can be added to other users groups without their consent.

Complaint 22 – New Policy
This relates to recent changes made to Facebook's Privacy Policy. The complainant contended that it is difficult to understand the changes in conjunction with the previous policy and that users have not had any opportunity to consent to the changes made.

Complaint against Facebook from the Norwegian Consumer Council

Background
In May 2010, the Norwegian Consumer Council lodged a complaint against Facebook with the Norwegian Data Protection Agency (Datatilsynet). The complaint was made on foot of concerns arising from an examination by the Council on the affects of social media use on consumers.

The Norwegian Data Protection Agency considered that Norwegian law did not apply in relation to the complaint and, given that Facebook Europe is located in Dublin, that Irish law should apply and that the matter should be addressed by the Irish authorities.

Complaint
The complaint is generally described as being in relation to the collecting, processing and storing of personal data on Facebook. In particular, the complaint examines Facebook's privacy settings and information passed to third party applications accessible from the Facebook platform.

Privacy Settings
The complainant highlights a number of changes made by Facebook to privacy settings functionality. In one instance in December 2009, the Council considers that the new privacy settings recommended by Facebook would allow certain information, for example 'posts by me' and 'religious views' to be available to a wider user audience and that "members were urged to accept the new privacy settings".

The Council also takes issue with another change, stating that, formerly, it was possible for a user to block all third party applications with a simple click, but now they had to be removed individually.

Third Party Applications
The Council highlights the fact that when a Facebook user signs up to a third party application that the user's data is provided to the application. The Council contends, from information collected from a survey it carried out, that "many people believe the applications to be part of Facebook and they are therefore not even aware that they are interacting with a third party". The Council also considers that many of the terms and conditions of third party applications are complex or unclear.

The Council states that a user signing up to a third party application must accept the application's terms and conditions in order to use the service. The Council contends that Zynga, a third party application providing 37 different games, collects and holds the following information when you sign up:

- Name
- Address
- Gender

The Council has concerns in relation to additional data subsequently collected:

- Information about you from other Zynga users
- Information on those you invite to use a Zynga game

The Consumer Council states that there is no information provided on

- How long this data is stored
- What it is used for
- Whether it is shared with others

The Council also contends that Zynga retains personal data even after a user account has been closed. The Council highlighted the use of this retention of data with an example from the Zynga

game Café World, providing a screenshot of the game which shows a profile photo of a Facebook user continuing to appear in the game after the user had previously blocked the application.

Conclusion

The Norwegian Consumer Council sets out a number of areas under which it considers Facebook and Zynga to be in breach of Norwegian Data Protection Law and raises issues which may be summarised by the following questions:

- Are Facebook and Zynga legally entitled to collect, store and process sensitive personal data such as race and political affiliation?
- What responsibility does Facebook take for third party applications offered via its platform?
- Given the complexity of third part terms and conditions, can a user be deemed to have provided free and informed consent to the processing of their data?
- Are Facebook and Zynga complying with data protection law in terms of limitations to the processing of personal data?
- Is the processing of personal data by Facebook and Zynga compatible with and relevant to the purpose for which it was collected?
- Is the data being retained for longer than is necessary, given the purpose for which it was originally collected?
- Is the disclosure of data by both Facebook and Zynga in breach of data protection law?
- The Council states that while Facebook is on the 'Safe Harbour' list, Zynga is not. Accordingly, does Facebook's processing of personal data in the context of its relationship with Zynga conflict 'Safe Harbour' rules?

Appendix 3

Overview of Team Functions

(Provided by FB-I)

Team	Main Description	High Level List of Activities	Region of Responsibility
Developer Relations	Developing the Facebook site and APIs. Providing worldwide technical support to third party developers who use our public APIs to develop Apps and Games for use on/with Facebook	Developing new features for Facebook site, particularly API code. Debugging Facebook Site and API Code Answering technical questions raised by the sales teams in regard to their clients' Facebook integrations Investigating and resolving issues with developers' accounts and their application settings/configuration Reviewing and addressing bug reports from other developers, reproducing the issue and triaging to Platform Engineering, if additional support required.	Global / EMEA (Sales client support is primarily EMEA, other functions global)
Site Reliability Operations	Team of Operations Engineers – providing front line management and support primarily for the core server infrastructure	Deploying new infrastructure Developing the tools used to manage the server fleet Maintaining existing infrastructure Handling system administration for the development & production environment Monitoring and responding to alarms Engaging with and escalating high severity site events	EMEA
Network Operations	Team of Operations Engineers – providing front line management and support primarily for the core network Infrastructure	Deploying new infrastructure Maintaining existing infrastructure Monitoring and responding to alarms Engaging with and escalating high severity site events	EMEA
Database Operations	Team of Operations Engineers – providing front line management and support primarily for core database infrastructure	Deploying new infrastructure Maintaining existing infrastructure Handling system administration of the development and production environment Monitoring and responding to alarms Engaging with and escalating high severity site events.	EMEA

Legal	Ensuring, with outside counsels' assistance, that all Facebook products comply with all applicable regulations, including data protection laws. Providing, with outside counsels' assistance, general legal support to FB teams in EMEA.	Working with Facebook's global engineering and legal staff and outside counsel to ensure that all Facebook products and policies are developed in accordance with European and Irish privacy laws. Ensuring all FB polices and product comply with all EU and national laws, e.g. advertising policies, risk payment policies, etc. Advising FB operational team on the implementation of FB policies. Drafting contractual documents for sales teams based in EMEA (OSO and ISO in Dublin and DSO in local offices); Advising support functions (human resources, finance, etc.) with regard to EU and national legislation and regulation Interfacing (in collaboration with policy team) with EU authorities (competition, privacy, consumer protection, etc.) Liaising with data subjects in relation to subject access requests	EMEA + India
Public Policy	Developing and communicating Facebook policies to regulators, legislators etc.	Working with legislators and regulators to explain Facebook policies (particularly privacy policies) and to resolve complaints. Communicating with the media about Facebook policies. Liaising with Legal and User Operations in relation to subject access requests and other privacy issues.	EMEA
Platform Operations	Enforcing our Platform Policy to protect our users. Engaging proactively and reactively with developers, sales team and partner teams with regard to the platform and in particular the platform policy.	Enforcing Platform Policy against applications and developers to protect user data and Facebook. This is done through both automated processes and managing appeals. Educating stakeholders on Platform and Platform policy. Consulting with stakeholders such as developers to ensure that applications properly implement policy and provide a good user experience. Investigating all reports concerning applications and taking enforcement action, if necessary.	EMEA and APAC

User Operations	Maintaining a safe environment for our users by enforcing our Statement of Rights and Responsibilities and Privacy Policy. Providing proactive and reactive assistance to users across the EMEA region.	Ensuring that Facebook platform remains a safe environment for our users and developing and implementing policies and technologies to achieve this objective. Enforcing the SRR by investigating allegations of inappropriate behavior on Facebook, and taking appropriate action where necessary. Ensuring compliance with rules prohibiting abusive content such as imposter accounts, bullying, extremism, defamation, impersonation, graphic or inappropriate content and IP infringement. Responding to subject access requests. Providing email assistance to our users, across EMEA, in 20 languages. Typical issues include addressing hacked accounts, determining appeals from banned users, and IP infringement. Proactively flagging and reviewing content in breach of our SRR to provide a safe environment for our users. Generating and providing user insights and user satisfaction reports for the region.	EMEA - 20 supported languages.
Risk Operations	Mitigating, on a worldwide basis, any financial loss or compliance breach found on our platform (Developer and Advertiser platforms) by proactively detecting and investigate potential fraud and taking action to stop it.	Daily monitoring of transactions hitting fraud rulesets. This is subdivided into Credits Buyer, Credits Seller and Advertiser traffic. Reviewing spending limits for all customers. Liaising with affected users.	Global
Payment Operations	Global responsibility for Facebook Credits. Ensuring payment performance to enable revenue and reduce costs. Working with and supporting our global user base using Facebook Credits.	Ensuring optimal user experience, with significant focus on our EMEA/APAC payments product Integrating new payment methods for Facebook Credits Working with developers to enable them to incorporate Facebook Credits into their applications so as to enable effective monetisation. Resolving user issues when purchasing Facebook Credits	Global

Online Sales Operations	Engaging proactively and reactively with OSO advertisers (managed and unmanaged) across the EMEA region.	Ads policy support for all channels and all advertisers in EMEA (DSO, ISO, OSO) Account management of a segment of ISO/OSO advertisers, including strategic support of advertisers Facebook Platform usage (media, Pages, Social Plugins) Account Specialist support for managed ISO and OSO clients (Ads & media, Page, and Social Plugin optimisation & strategy; "best practice" collateral creation) Advertiser analysis and reporting SMB support and education for higher-value unmanaged advertisers through client-facing and scalable initiatives (partnering with SMB marketing) Management of reactive and automated support for unmanaged OSO advertisers	EMEA, some LATAM and APAC (but moving these to the relevant regions)
Inside Sales Operations	Driving ISO prospective efforts across the EMEA region. Engaging proactively and reactively with ISO accounts (direct clients and agencies) across the EMEA region.	Evangelising and driving strategy around Facebook Platform (ads, apps, pages, social plugins etc…) Acquiring and managing direct accounts and agencies Managing the inbound leads queue Managing outsourced lead generation programme in EMEA Acquiring lead clients using several prospecting sources Developing RFPs for Premium and Marketplace ads Representing Facebook at speaking opportunities at various key industry events.	EMEA and LATAM
Advertising Operations	Managing Premium Advertising campaigns for the EMEA region Managing all advertising operations for EMEA DSO clients (from Q1 2012)	Building and managing performance of premium advertising campaigns Working directly with DSO/ISO clients to collect advertising assets Developing ad performance strategies, contributing to overall DSO sales efforts Working with Revenue Assurance to ensure that all activity can be correctly invoiced Working directly with DSO clients to ensure optimal performance of Marketplace campaigns (from Q1 2012)	EMEA

Marketing	Driving acquisition, growth and retention of EMEA based small and medium advertisers	Co-coordinating Facebook Pages Strategy for EMEA Developing and implementing onsite marketing strategy for EMEA Developing and implementing Facebook Ads and SEM strategy for EMEA Generating leads for ISO and OSO call centres Running acquisition and retention webinar events and programmes and email marketing acquisition and retention programmes.	EMEA & occasionally Global depending on project
Finance	Finance functions serving most business needs of all FB offices outside North America	Order to Cash Functions; assessing customer credit worthiness, reviewing FB IOs for revenue compliance, billing, collections-Procure to pay cycle - vendor management, logistics, accounts payable-Record to Report - monthly reporting by FB company excl North America & consolidation into US parent-Compliance - tax, statutory returns & financial statements, statutory audits-Payroll, tax planning, treasury & banking functions-Business Operations - partner with ad sales & user centric teams on strategy, prioritization, system enhancements, performance reporting, sales compensation programs, resource planning	Mostly Global excluding North America
Learning & Development	Providing learning and development opportunities to all staff. We provide both job specific and transitional knowledge and skills	Departmental/organisational training needs analysis Collection of individual training needs Development of required programmes Delivery of required programmes Evaluation of programmes 3rd party vendor management	EMEA
Human Resources	Providing full HR service for employees and managers within the EMEA business for all stages of the employee life cycle. Supporting the delivery of companywide HR initiatives such as Performance Management,	New hires: Accountable for new hires receiving relevant starter documentation, support orientation programme, set up new hires on HR systems and payroll. Day to Day administration: Providing administrative support to employees with Day to Day queries, requests, updating employee files (e.g. landlord reference letters, visa application) Partner business: Supporting managers in business area's with business goal delivery related to people, coach and	EMEA

	Employee Engagement. Responsible for the design, roll out and adherence to HR policies and practices ensuring compliance within region.	mentor our people managers through the employee life cycle. Engagement: Driving initiatives within the business to enhance the Engagement of all employees. Performance management: Facilitate the company wide performance management process bi-annually and supporting managers with low performance issues. Policies & Procedures: Reviewing, amending and apply companywide HR policies and Procedures, ensuring compliance levels met. Transitions: Manage employee transitions within and from the business.	
Staffing	Hiring of employees for EMEA from EU HQ in Dublin	Hiring staff across regions, roll out HR programs, screening CV's & interviews, driving offer process, supporting all client groups & maintaining data reports and performance metrics on a regular basis.	EMEA
Real Estate & Facilities	Providing Real Estate & Facilities Support for the various offices in EMEA.	Managing and Developing EMEA real estate portfolio. Running Facility Management function to develop most suitable working environment for EMEA employees by applying best practice to improve efficiency by reducing operating costs while increasing productivity. Supporting other business functions in both strategic planning and data-to-day operations.	EMEA/ APAC
Physical Security	Physical Security Support to all teams in our offices within EMEA	Development of security procedures and policies for EMEA offices Providing physical security, including access control, to our EMEA offices. Physical security guarding and security mobile response function. Safety management programme for traveling employees Conducting security risk assessment and assisting with security systems designs for offices. Assisting Real Estate/ Facilities with our safety programmes and initiatives at EMEA offices.	

Appendix 4

Structure of European Offices

(Supplied by FB-I)

1. Introduction

FB-I Limited, based in Hanover Quay, Dublin, offers the Facebook platform to all users outside of North America.

With approximately 400 staff, Hanover Quay is Facebook's largest office outside of Silicon Valley and the group's EMEA headquarters. It is the only office, and legal entity, within the Facebook group with control over non-North American user data.

As part of the expansion of Facebook's advertising business, a number of regional offices have been opened across Europe. These offices, which are under the management of FB-I, have been established to more effectively build links with local advertisers and, in some cases, to support the local developer community. These regional offices, are, in effect, local sales and support offices. While these local offices sell advertising on the Facebook platform, they are not involved in the development or governance of the site or the control of user data. A limited number of staff in these offices have limited access to user data so as to generate reports about the effectiveness of advertising and to troubleshoot potential issues for the advertisers with which they deal. Such access is highly regulated pursuant to data processing agreements entered into between FB-I and the various European subsidiaries copies of which are being provided to the DPC.

2. Role of FB-I

FB-I's 400 staff (around 326 Full Time Employees and 75 contractors) are responsible for the development and maintenance of the Facebook platform, the protection of Facebook users, the corporate administration of many of Facebook's non-North American activities and the sale of advertising to customers.

2.1 Engineering and Platform Development

FB-I's engineering staff are divided into two teams: Developer Relations and Infrastructure Engineering. The Developer Relations team are responsible for the API's which app developers use to interact with the Facebook platform and the personal data which Facebook users have chosen to share.

The Infrastructure Engineering team manages site reliability, network infrastructure and databases. This team deploys new infrastructure as required and ensures that Facebook remains functional.

Aligned to the engineering teams, FB-I's Payment Operations group has global responsibility for the development and rollout of Facebook Credits, the currency of the Facebook platform.

2.2 Legal and Policy

FB-I operates in accordance with Irish law. Irish data protection law regulates the use of all non-North American Facebook user data. Ensuring legal compliance is the role of the Legal team. The Dublin based Legal team works with Facebook's global engineering and legal staff and outside counsel to ensure that all Facebook products and practices comply with all applicable regulations, including data protection laws.

The Policy team is responsible for Facebook's public reputation and works closely with the media, legislators and regulators to explain Facebook's practices and policies to decision makers and the public at large and to take feedback from these stakeholders.

2.3 User Safety

FB-I has over 100 staff dedicated to the monitoring of the Facebook platform and the protection of Facebook users. These staff are spread across the User Operations, Law Enforcement Response, Risk an Payment Operations, Developers Relations and Platform Operations Teams.

Allegations of unlawful conduct by Facebook users are investigated by the User Operations team, which has responsibility for enforcing the Facebook Statement of Rights and Responsibilities and Privacy Policy. This team, which operates in over 20 languages, is charged with ensuring that Facebook remains a safe environment. The team investigates allegations of inappropriate behaviour on Facebook and takes action against abusive content such as imposter accounts, bullying, extremism, defamation, impersonation, graphic or inappropriate content and IP infringement.

The Law Enforcement Response Team is responsible for reviewing requests from law enforcement authorities for access to Facebook user data. This team reviews all received requests against the Facebook Privacy Policy and applicable law to ascertain if there are legal grounds justifying the handover of such data.

The Risk Operations team is responsible for protecting, on a worldwide basis, Facebook users from potential fraud on the platform.

The Platform Operations team polices applications to prevent any misuse or abuse of user data. This team enforces the Platform Policy and takes action against any developers in breach. This team also has an educational function, teaching app developers about their responsibilities to users and the rules governing the Facebook platform.

2.4 Managerial Functions

As EMEA headquarters, FB-I is responsible for managing regional European offices. This regional management is the responsibility of a number of executives and supporting teams based in Facebook's Dublin office, including Finance, Recruitment/HR, Real Estate/Facilities and Security.

The twenty-three person finance team handles the finance functions for all Facebook offices outside of North America.

Facebook's non-North American recruitment and training is managed by FB-I. Twenty-four recruitment specialists are presently based in Hanover Quay and are responsible for sourcing the global talent underpinning Facebook's continued expansion. Non-North American staff, once recruited, are trained primarily in the Dublin Office. In rare circumstances, certain specialist staff members might also receive training in Facebook's Silicon Valley offices. Only one Dublin staff member received their initial training in Palo Alto in 2011. Human resources administration for EMEA staff is handled by FB-I.

The development of new Facebook offices in EMEA and Asia Pacific is the role of the Real Estate/Facilities team. This team identifies potential new properties which could be acquired to facilitate the continued growth of the business. This team also ensures that all offices provide a suitable working environment for Facebook staff.

The safety of Facebook staff and the security of Facebook offices is a paramount concern. The physical security team manages the security of both Facebook's Dublin HQ and the regional offices.

2.5 Advertising and Sales

Facebook earns revenues through the sale of advertisements on its platform. FB-I employs nearly 150 people to develop and sell advertising products across the EMEA region. These teams, which are divided up based upon the size of potential advertising customers, work with our advertising partners to promote the benefits of social advertising on Facebook and to ensure that such advertising customers make the best use possible of Facebook. These teams also work with the various regional Facebook offices to support the local advertising market.

3. Regional Offices

FB-I is supported by a network of small local offices scattered across the EU. These offices, which operate under the control and direction of FB-I, seek to promote Facebook, and Facebook advertisements, within their geographical and linguistic area of responsibility.

These offices have no role in the development or maintenance of the platform or the control of user data. Their functions are limited to the sale of advertising, local PR and, in limited cases, addressing queries from local app developers. In the context of carrying out these duties, these offices may process a limited amount of user data relating to the pages of advertisers and prospective advertisers pursuant to processing agreements entered into with FB-I.

Offices are located in Amsterdam, Hamburg, London, Madrid, Milan, Paris, Stockholm.

Appendix 5

Law Enforcement Requests

(Provided by FB-I)

1. Introduction

FB-I Limited, as the data controller for non-US and Canadian user data, receives and processes requests for user data from law enforcement agencies around the world. The following memorandum describes the law and policies that guide the processing of these requests.

2. Policy Framework: Introduction

We disclose account records solely in accordance with our terms of service and applicable law. Our Privacy Policy, to which all users must consent before using our service, sets forth our policies relating to the disclosure of user information in response to legal requests:

> *We may share your information in response to a legal request (like a search warrant, court order or subpoena) if we have a good faith belief that the law requires us to do so. This may include responding to legal requests from jurisdictions outside of the United States where we have a good faith belief that the response is required by law in that jurisdiction, affects users in that jurisdiction, and is consistent with internationally recognized standards. We may also share information when we have a good faith belief it is necessary to: detect, prevent and address fraud and other illegal activity; to protect ourselves and you from violations of our Statement of Rights and Responsibilities; and to prevent death or imminent bodily harm.*

Facebook has adopted a tiered approach to the disclosure of user data which draws a distinction between user-generated content data (such as messages, photos, and videos) and non-content data (such as basic subscriber information and logs of IP addresses). This framework reflects similar classifications set forth in a U.S. privacy law, the Stored Communications Act, which provides robust protections for electronically stored information.

Adopted first by Facebook, Inc., FB-I has continued to respect this distinction between content and non-content data, as it comports with law; affords strict protections to user data; provides law enforcement with avenues to obtain data in criminal investigations pursuant to law and necessity; and ensures proportionality in our responses to law enforcement requests. In many cases, for example, law enforcement agencies seek only information which may help them locate a potential suspect or missing person. The disclosure of basic subscriber and IP address information, rather than account content, routinely suffices for such purposes.

Our approach to law enforcement requests can be broken down into the following categories: (1) non-content requests; (2) content requests; and (3) emergency requests.

2.1 Non-Content Requests

FB-I scrutinises requests for non-content information from law enforcement under the three prongs of the Privacy Policy set forth above. If, at the end of this assessment, we are satisfied that the request meets each of these requirements, we may disclose *basic subscriber information* to the requesting agency.

a) *Does the request meet legal standards in the originating jurisdiction?*
Under the guidance of legal counsel, we review the validity of the request in light of the applicable local laws, which vary by jurisdiction. Where we are uncertain as to the validity of the request

under local law, or the identity of the requesting body, we often directly consult local counsel in the relevant jurisdiction.

b) *Is the request intended to protect Facebook or our users in that jurisdiction?*
As a further limit to overreaching by local law enforcement officials, Facebook also evaluates requests to ensure that they affect Facebook or its users in the requesting jurisdiction.

c) *Is the request consistent with internationally recognised standards?*
We will decline to disclose user data in response to law enforcement requests that are inconsistent with internationally recognised standards, such as freedom of speech.

We may disclose basic subscriber information if a requests meets these requirements. We do not, however, disclose content information, in response to a request for non-content data.

2.2 Content Requests
Should the law enforcement agency require content information from FB-I, we will require that we be served with a legally compelling request under Irish law. The Gardaí will be required to produce a search warrant or similar coercive document. Non-Irish search warrants will only be respected by FB-I if they are enforceable as a matter of Irish law. This will require that any such orders be domesticated by way of application to the Department of Justice pursuant to the Criminal Justice (Mutual Assistance) Act 2008.

2.3 Emergency Requests
Consistent with our Privacy Policy, we may disclose user data where we have a good faith belief that disclosure is necessary to "*prevent death or imminent bodily harm.*"

3. Implementation of Policy and Process

The above policy is implemented by the FB-I Law Enforcement Response Team. This four-person team (which will continue to grow) is based in our Dublin Office and has responsibility for law enforcement requests excluding those from North America. They are complemented by a team based in Palo Alto which focuses on the Americas. These teams support one another and their respective operations to ensure that Facebook can provide a 24/7 service to law enforcement.

3.1 Compliance with the Acts
The manner in which we disclose personal data to law enforcement is full in compliance with the requirements of the Data Protection Acts 1988 and 2003.

3.2 Basic User Information Disclosure
The disclosure of basic user information to law enforcement authorities, in accordance with our Privacy Policy and applicable law, is consistent with the Acts.

First, our limited disclosures of user data in response to these official requests comport with the Acts because they fall within the law enforcement exemption. In particular, Section 8(b) provides that the restrictions on processing laid down in the Acts do not apply if the processing is *"required for the purpose of preventing, detecting or investigating offences, apprehending or prosecuting offenders or assessing or collection any tax, duty or other money's owned or payable to the State, a local authority of a health board, in any case in which the application of those restrictions would*

be likely to prejudice any of the matters aforesaid." This provision exempts any assistance provided to the Gardaí and other (non-Irish) law enforcement authorities from the requirements of the Acts, to the extent that such processing was *necessary* to provide such assistance and would be prejudiced by the application of the Acts.

Fundamental principles of statutory construction support the conclusion that non-Irish law enforcement requests (and especially those of other member states of the EU) fall within the Section 8 exemption.

The doctrine of *inclusion unis exclusion alterius* states that where a statute expressly lays down a number of restrictions, it is to be assumed that omitted restrictions were deliberately omitted.[52] A number of the exemptions in Section 8 are subject to express territorial limits. Section 8 refers, in places, to "the security of the State", "monies payable to the State" and "the international relations of the State". In contrast, the exemption in Section 8(b) for investigating and prosecuting offences is not qualified by the words "within the State." Applying the *inclusion unis exclusion alterius* canon, we can infer that this territorial restriction was deliberately omitted by the Oireachtas.

In addition, this reading would be in keeping with the text of the Directive, to which regard must be paid when interpreting the meaning of the Acts.[53] Section 8 of the Acts must be read in light of Article 3(2) of the Directive. This provides that the Directive does not apply to the processing of personal data "*in the course of an activity which falls outside the scope of Community law, such as those provided for by … Title VI of the Treaty on European Union [criminal law] and in any case to processing operations concerning public security defence, State security (including the economic well-being of the State when the processing operation relates to State security matters) and the activity of the State in areas of criminal law*." Given that "State" in the context of the Directive refers to the member states of the European Union, this strongly suggests that the exemptions for criminal investigations in Section 8 extend to investigations outside the state, at least to the extent that such investigations are being conducted by other member states of the EU. This reading is further supported by Recital 9 to the Directive which makes clear that Member States cannot apply different data protection standards to processing inside the state and processing with a pan-EU dimension:

> *Given the equivalent protection resulting from the approximation of national laws, the Member States will no longer be able to inhibit the free movement between them of personal data on grounds relating to protection of the rights and freedoms of individuals, and in particular the right to privacy.*

Finally, while there is a general presumption that criminal legislation is not to have extra-territorial effect,[54] the Courts have found that civil law statutes which are silent on the question can cover

[52] *Kiely v. Minister for Social Welfare* [1977] IR 267; *Fanning v. University College Cork* [2007] 18 ELR 301; *Inspector of Taxes v. Arida* [1992] 2 IR 155.
[53] *C-106/89 Marleasing* [1990] ECR 4135; *Lawlor v. Minister for Agriculture* [1990] 1 IR 356; *Bosphorus Hava v. Minister for Transport* [1994] 2 ILRM 551; *Murphy v. Bord Telecom Éireann* [1989] ILRM 53; *Campbell v. MGN* [2003] 2 WLR 80.
[54] *DPP (Broderick) v. Flanagan* [1979] IR 265.

issues arising out of jurisdiction. In *Murphy v. GM*[55] Justice O'Higgins, finding that a receiver appointed under the Proceeds of Crime Act 1996[56] could use his powers against assets outside the jurisdiction (the Act being silent on this issue) held that *"there is nothing in the Act that requires the reading down of the words so as to imply that the Act applies to assets within the jurisdiction only. If the powers of the receiver were to be limited territorially the legislation would have so stipulated. Nor is there any legal principle or principle of interpretation that requires such a narrowing down. I cannot agree that there are any internal indications in the Act to suggest intra-territorial application only.*"

Accordingly, Section 8(b) applies to investigations occurring outside of Ireland, and our disclosures to law enforcement authorities therefore fall outside the Acts.

Even if it were assumed that a Section 8 exemption did not apply, however, the limited and careful processing of user data in response to law enforcement requests is nevertheless justified and lawful under Section 2A of the Acts. Our users consent to such disclosure through the Privacy Policy. Furthermore, even in the absence of such consent, such disclosure is lawful and appropriate in light of Facebook's legitimate interest in cooperating with law enforcement authorities.

As noted above, our Privacy Policy provides:

> *[We may share your personal information in] responding to legal requests from jurisdictions outside of the United States where we have a good faith belief that the response is required by law in that jurisdiction, affects users in that jurisdiction, and is consistent with internationally recognized standards.*

Our users choose to use Facebook in light of this commitment, which specifies the circumstances under which disclosures may be made to law enforcement authorities. Furthermore, Section 2A(1)(d) of the Acts also supports our limited assistance to law enforcement authorities. This provision allows for processing where necessary for the purposes of the legitimate interests of either Facebook or the third party to whom the disclosure is made – in this case, law enforcement officials. As a matter of Irish law, a business has a legitimate interest in cooperating with law enforcement.[57] Law enforcement authorities, of course, have an official duty to properly seek out evidence and investigate offences.[58]

We have adopted effective procedures to ensure that we only disclose user information in accordance with law and our Privacy Policy. As noted above, we disclose only a limited amount of user data and law enforcement cannot require the disclosure of user content in the absence of a coercive warrant or order. Furthermore, and in accordance with Section 2A(1)(d), we scrutinise

---

[55] *Murphy v. GM* [1999] IEHC 5

[56] Notwithstanding its name, this statute is a civil, and not a criminal, law measure: *Gilligan v. Criminal Assets Bureau* [1998] 3 IR 185.

[57] *DPP v. Forbes* [1993] ILRM 817; *Minister for Justice v. Wang Zhu Jie* [1991] ILRM 823.

[58] *Dillon v. O'Brien and Davis* (1887) 20 LR IR 300; *Braddish v. DPP* [2001] 3 IR 127; *Dunne v. DPP* [2002] 2 IR 305; *Scully v. DPP* [2005] 1 IR 242; *O'Callaghan v. Judges of Dublin Metropolitan District Court* [2004] 2 IR 442; *Ludlow v. DPP* [2009] 1 IR 640.

law enforcement requests to ensure that law enforcement is not overreaching or seeking data in matters that would infringe upon basic freedoms of speech.

3.3 Content Information

As noted above, we will only provide user content information where we are compelled to do so under Irish law. Such an obligation can either result from a warrant or other statutory power of compulsion (on the part of the Gardaí) or as a result of some non-Irish order which has been domesticated through the process laid down by the Criminal Justice (Mutual Assistance) Act 2008. In either case, we are obliged to accede to the terms of the order. Consequently, any processing in which is necessary by such compliance is exempted from the Acts by virtue of Section 8(e) of the Acts which exempts processing which is "required by or under any enactment or by a rule of law or order of a court".

3.4 Emergency Information

Disclosure to law enforcement authorities in emergencies is permissible in light of Section 8(d) of the Acts, which exempts processing which is required urgently to prevent injury or other damage to the health of a person or a serious loss of or damage to property. We will only disclose user data if required to avoid *death* or *serious personal injury.*

The disclosure of user data to prevent harm in these circumstances is not only compatible with the Acts, but is also in line with the priority Irish law gives to the preservation of life. In *People (DPP) v. Shaw,*[59] it was accepted that the constitutional right to life takes precedence over other constitutional rights. In that case, the need to protect a life in imminent peril justified the police detaining the suspect after the elapse of the statutory detention period: the right to life took precedence over the right to liberty. The Acts must be interpreted in light of the precedence the Irish Constitution affords to right to life.[60]

4. Data Export

Responding to law enforcement requests may involve the transfer of user data outside of the EEA. Such transfers comply with the requirements of Section 11 of the Acts based on the consent of our users and, in certain circumstances, as a requirement under an enactment. As noted above, we specifically inform our users, in our Privacy Policy, that "*[we may share your personal information in] responding to legal requests from jurisdictions outside of the United States."*

Furthermore, in cases where we are compelled to disclose user data on foot of an order made pursuant to the Criminal Justice (Mutual Assistance) Act 2008, the export of such user data outside of the EEA is rendered lawful by virtue of Section 11(4)(a)(i)(I) which provides that the restriction on transfers does not apply if the transfer is required or authorised under an enactment.

5. Conclusion

FB-I's restrictive and responsible policies in responding to requests by law enforcement authorities for user data properly balances our duty to the public to respond to legitimate law enforcement enquiries while fully protecting the privacy of our users. The policy we have adopted is fully compliant with the requirements of the Data Protection Acts 1988 and 2003.

---

[59] *People (DPP) v. Shaw* [1982] 1 IR 1.

[60] *East Donegal Co-Operative Livestock Marts v. Attorney General* [1970] IR 317.

Appendix 6

Minors

Minors

Facebook in its **Data Use Policy**[61] (Privacy Policy) states

> **V. Minors and safety**
> We take safety issues very seriously, especially with children, and we encourage parents to teach their children about safe internet practices. To learn more, visit our Safety Center.
>
> To protect minors, we may put special safeguards in place (such as placing restrictions on the ability of adults to share and connect with them), recognizing this may provide minors a more limited experience on Facebook.

see https://www.facebook.com/safety

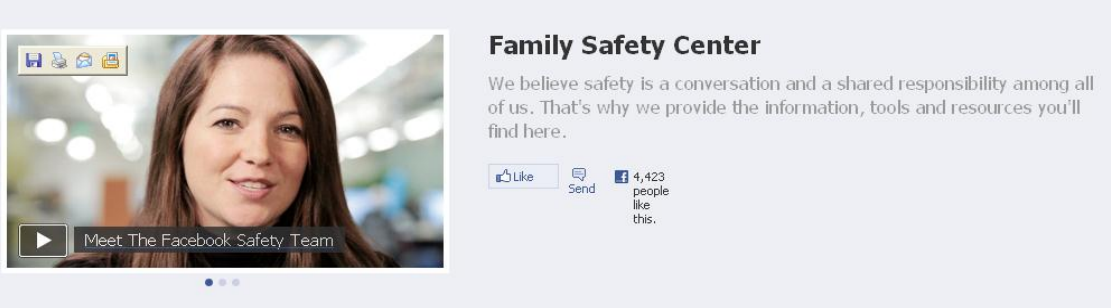

There are a wide array of tools and resources available to visitors broken down into their respective audiences (see next screen grab)

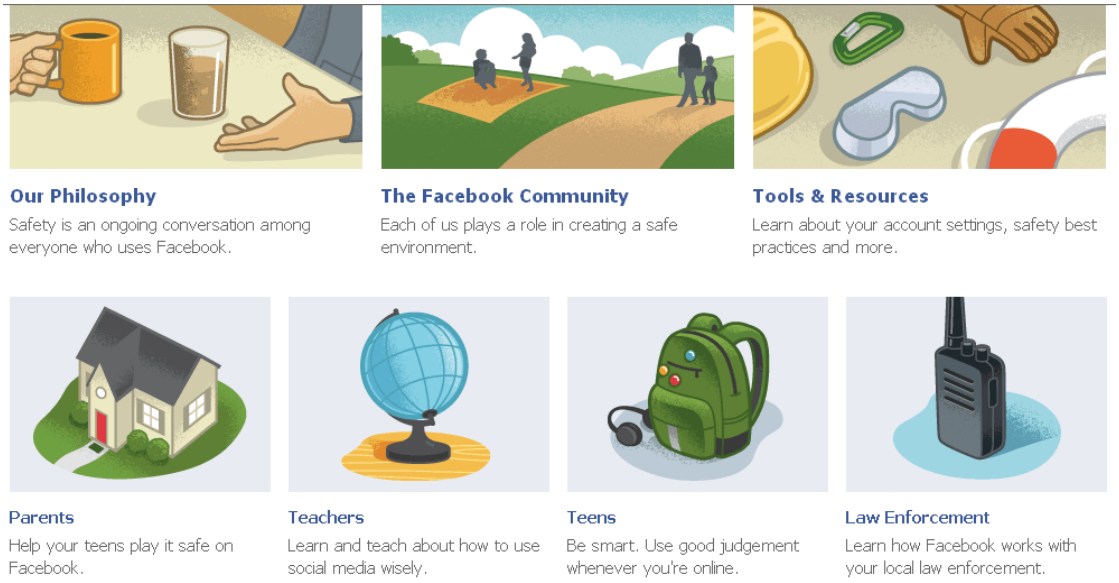

Clicking the icon for teens brings the next screen:

[61] http://www.facebook.com/full_data_use_policy

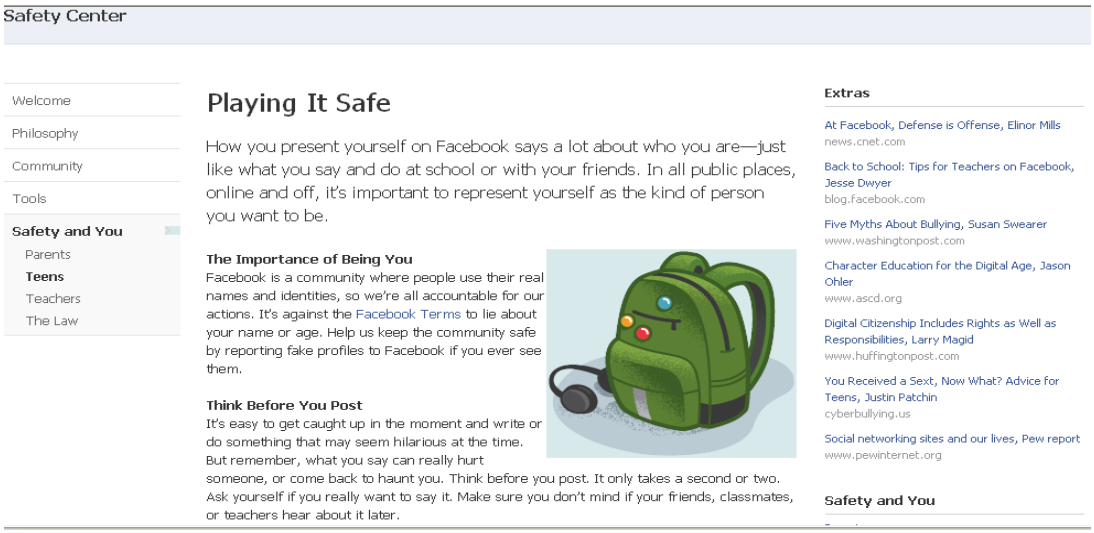

With the following information

Playing It Safe
How you present yourself on Facebook says a lot about who you are—just like what you say and do at school or with your friends. In all public places, online and off, it's important to represent yourself as the kind of person you want to be.

The Importance of Being You
Facebook is a community where people use their real names and identities, so we're all accountable for our actions. It's against the Facebook Terms to lie about your name or age. Help us keep the community safe by reporting fake profiles to Facebook if you ever see them.

Think Before You Post
It's easy to get caught up in the moment and write or do something that may seem hilarious at the time. But remember, what you say can really hurt someone, or come back to haunt you. Think before you post. It only takes a second or two. Ask yourself if you really want to say it. Make sure you don't mind if your friends, classmates, or teachers hear about it later.

Also remember that any information you post – whether in a comment, a note, or a video chat – might be copied, pasted, and distributed in ways that you didn't intend. Before you post, ask yourself - would I be OK if this content was shared widely at school or with my future employer?

At the same time, we all make mistakes. If you find yourself wishing you hadn't said or done something, it's never too late to apologize.

Don't Talk to Me Anymore
If you ever receive hurtful or abusive messages or posts on your profile page you have options. Depending on how serious the situation is, you can ignore it, ask the person to stop, unfriend or block the person, or tell your parents, a teacher, a counselor, or another adult you trust. Everyone deserves to be treated with respect.

Report Abusive Content

Be sure to always report abusive content—whether it's on your profile page, or someone else's. You can also report inappropriate Pages, Groups, Events and fake or impostor profiles. (Remember that reporting is confidential, so no one will know who made the report.)

Tips for Teens
Don't share your password with anyone.
Only accept friend requests from people you know.
Don't post anything you wouldn't want your parents, teachers, or employer to see.
Be authentic. The real you is better than anything you might pretend to be.
Learn about privacy settings, and review them often.

Functionality & Features on Facebook for minors.

Facebook provided a detailed overview of the differences between all of its privacy settings and sharing options for minors and adults.

How do minors connect with other people?

Like adults, minors can appear in search results. **A search cannot be conducted using age or location.** Adults and minors can also receive a friend request from someone who is not already a friend of a friend, such as a family relative or a friend with whom they have no mutual friends. Because friend requests may come from adults they don't know, minors should always be careful when accepting these requests.

Messages are handled differently for minors and adults:

Minors	Adults
Depending on their settings, minors may receive messages from people who are **friends of friends** on Facebook, which may include adults they don't know.	Depending on their settings, adults can be messaged by anyone on Facebook.

Facebook indicated to the Office that it encourages all users to visit their Privacy Settings page to manage and limit who is allowed to send you friend requests and messages.

How do minors share posts like photos and status updates?

Minors	Adults
Minors can share with a maximum of **friends of their friends.**	Adults can share posts with a maximum of everyone by posting as **public**.

How does tagging work for minors?

You can add tags of anyone — including people you don't know, any place, or any page (like a celebrity or sports team). When you add a tag of someone to your photo, status update or other posts, you may also be sharing the post with that person's friends. Keep in mind that the tagged person's friends might include adults you don't know.

Who can tag you?

Minors	Adults
Only friends of your friends can tag you in a post.	Any person can add a tag of you to their post

Who can see tags of you?

Visibility of tags of Minors is different for young people than adults. Visibility varies based on who added the tag and who is viewing the post.

Minors	Adults
If a **friend** adds a tag of you, that tag is visible to **any person** who can see the post on Facebook and in third-party applications like quizzes and games. If a **friend of a friend** tags you, only your friends can see your name with a link to your profile (timeline). This visibility applies to Facebook and third-party applications like quizzes and games. If a **friend of a friend** adds a tag of you, non-**friends** will just see your name with no link to your profile (timeline).	If **any person** adds a tag of an adult, the tag has the same visibility as the post itself on Facebook and in third-party applications like quizzes and games.

Who can add tags to your posts?

Any person who can see your posts may be able to add tags to them. Minors can share posts with a maximum of friends of their friends. We limit this when you're younger to help protect younger people's privacy.

Adding a tag of someone will share the post with that person's friends. You have the option to approve tags other people add to your posts before they appear, using a new Tag Review feature. To turn the feature "on" or "off," go to the How Tags Work setting in your privacy settings.

When you approve a tag, keep in mind that you are giving permission for the tagged person's friends to see your post, which might include adults that you do not know.

Minors	Adults
For minors, the Tag Review setting is turned 'On' by default. You can turn this 'Off' if you want to allow friends to add tags to your posts without reviewing the tags first. You	For most adults, this setting is turned 'Off' by default so you don't have to approve tags that friends add to your posts. You can turn this 'On' if you prefer to approve tags added to your content before they go on Facebook.

will always have to review tags from friends of friends.	

How do location services work for minors?

The sharing tool gives you the option to share where you are when you post. You can add location to photos and status updates. There are special safeguards for Minors:

Minors	Adults
When you first use the new sharing tool location feature is 'Off' by default. To turn it 'On,' simply click the button to add location to your post. This setting will continue to stay 'On' until you change it.	For most adults, when you first use the new sharing tool the location feature is 'On' by default. To turn it 'Off,' simply click the 'X.'

Who can tag you?

Minors	Adults
Only friends of your friends can tag you in a post.	Any person can add a tag of you to their post

Who can see tags of you?

Visibility of tags of Minors is different for young people than adults. Visibility varies based on who added the tag and who is viewing the post.

Minors	Adults
If a **friend** adds a tag of you, that tag is visible to **any person** who can see the post on Facebook and in third-party applications like quizzes and games. If a **friend of a friend** tags you, only your **friends** can see your name with a link to your profile (timeline). This visibility applies to Facebook and third-party applications like quizzes and games. If a **friend of a friend** adds a tag of you, **non-friends** will just see your name with no link to your profile (timeline).	If **any person** adds a tag of an adult, the tag has the same visibility as the post itself on Facebook and in third-party applications like quizzes and games.

Who can add tags your posts?

Any person who can see your posts may be able to add tags to them. Minors can share posts with a maximum of friends of their friends. We limit this when you're younger to help protect younger

people's privacy.

Adding a tag of someone may share the post with that person's friends. You have the option to approve tags other people add to your posts before they appear, using a new Tag Review feature. To turn the feature "on" or "off," go to the How Tags Work setting in your privacy settings.

When you approve a tag, keep in mind that you may be giving permission for the tagged person's friends to see your post, which might include adults that you do not know.

Minors	Adults
For minors, the Tag Review setting is turned 'On by default. You can turn this 'Off' if you want to allow friends to add tags to your posts without reviewing the tags first. You will always have to review tags from friends of friends.	For most adults, this setting is turned 'Off' by default so you don't have to approve tags that friends add to your posts. You can turn this 'On' if you prefer to approve tags added to your content before they go on Facebook.

How do privacy controls change when a minor becomes an adult?

When you reach a legal adulthood, we will notify you that you are now considered an adult on Facebook and remind you of the new privacy controls and sharing options available to you as an adult.

Here is what you will be notified of upon adulthood:

1. You can now communicate with anyone on Facebook
2. Any person can add a tag of you to their posts
3. If someone adds a tag of you in a post and shares with Public, then that tag of you will be visible to everyone
4. You have the option to be contacted by everyone via the Public setting

How does the Public setting work for minors?

Minors can share with a maximum of friends of their friends. So, wherever minors see Public as an audience option, it means **friends of friends**. In contrast, when adults choose Public their audience selection includes **everyone on the internet**.

Why do minors see a Public sharing option for their profile (timeline) information?

Wherever minors see Public as an audience option, it means **friends of friends**. Once a minor reaches legal adulthood, the meaning for Public changes to **everyone on the internet**.

Example: If a minor chooses to share his hometown with Public, his hometown only displays to friends of friends. No one can be searched by location. When he reaches legal adulthood, his hometown becomes available to everyone on the internet.