Social network security

A SysSec Whitepaper*

September 4, 2012

1 Introduction

In recent years, social networks have become more than a technology. They directly influence the lives of millions of people around the world. Friendships, social interaction and shared media are just a small subset of the offered functionality. However, the growing popularity also comes with a downside. With over 800 million users [5] in December 2011, Facebook is the largest, most widely accepted social network so far. Recently, it was repeatedly referred to as being the Microsoft Windows of the smartphones. The large amount of information published, and often publicly shared, by users on their online social network profiles is additionally attracting the attention of attackers. If just a single successful attack is launched against a network such as Facebook, the impact is tremendous with over 800 million people being potential victims. To make sure that such an attack does not happen on a large scale, security researchers focus on various properties of these virtual communities and try to find solutions for arising problems.

Naturally, *pure* social networks like Facebook and its predecessors are very good examples and can be used as a reference for most case studies. There are, however, various other platforms to consider. A good example are gaming platforms like Steam [12], Origin [13] or BattleNet [11] where users interact, share their latest achievements or simply chat with each other. Other networks such as LinkedIn or Xing focus on more professional participants to help them establish business relationships and maintain them. In fact, a lot of communities reaching from the aforementioned gaming to research communities, already established their own social network to help likeminded individuals to keep in touch.

What all of these platforms have in common is the fact that they rely on their user's social interactions to function. They only differ in the validity of the presented persona and, from an attacker's point of view, the asset connected with the person behind that persona. That can be a real name and personal

*The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no 257007.

information on Facebook, credit card information on gaming platforms or in-game currency in an MMOG. Security researchers aim to protect those assets by devising new protection mechanisms or identifying previously unseen threats. This task is not always simple and, due to the unpredictable nature of humans and their actions, often challenging.

2 Traditional attacks

Attacks on social networks are usually variants of traditional security threats (such as malware, worms, spam [15], and phishing [14]). These "common" threats are thoroughly discussed in existing research papers. The one thing these attacks have in common when used in junction with social networks is their possibility to leverage personal data for a higher impact. Spam, for example, can be directly sent to an interested person, probably with the name of a friend as the sender [15]. Worms and other malware have a higher infection rate because links within a social network are more likely to be clicked [10]. Phishing attacks can be aimed at a narrow category of individuals with a higher success rate as traditional spam [4]. These attacks are carried out in a different context by leveraging the social networks as a new medium to reach the victims. Moreover, adversaries can take advantage of the trust relationships between "friends" in social networks to craft more convincing attacks by exploiting personal information gleaned from victims' pages. Therefore, most of the attack requires, as a first step, to become friend of the victim. As already mentioned in the introduction, that applies to almost any form of social networks as long as they support some form of "friendship".

As web applications served to the user via standardized, well-known protocols, social networks can also be attacked in equally well-known ways. OWASP lists the top ten of the web vulnerabilities which of course also apply to social networks. Placed on the very top are injection vulnerabilities. One might think that textbook-like SQL-injection attacks are a thing of the past, but in May 2011, they were the reason for roughly 56.000 user credentials of the dating-social-network findfriendz.com being disclosed. Facebook itself has been shown to be vulnerable to XSS (Cross-Site-Scripting) and CSRF (Cross-Site-Request-Forgery) attacks in the past [2].

While traditional attacks undoubtedly have a severe impact on the customer base provided by today's social networks, new attack vectors, which are specifically tailored to operate on the unique structure of social networks, are emerging.

3 New attack vectors

As the name already suggests, social human interaction is an integral part of social networks. Hence the user itself, rather than the technical infrastructure, is predominantly targeted by social engineering attacks. A typical example is to

spike a user's interest on a certain topic that in turn provokes an inconsiderate user action (scamming). A good example for this behavior are various "viral videos" that spread through Facebook over the last year. The new aspect in social engineering attacks in social networks are the trust relationships built upon the aforementioned "friendships". In fact, past research has shown that users of online social networks tend to exhibit a higher degree of trust in friend requests and messages sent by other users (e.g., [7, 9]).

In a *reverse social engineering* attack, this heightened amount of trust is exploited by an attacker that does not initiate contact with the victim. Rather, the victim is tricked into contacting the attacker herself. As a result, a high degree of trust is established between the victim and the attacker as the victim is the entity that first wanted to establish a relationship. Once a reverse social engineering attack is successful (i.e., the attacker has established a friend relationship with the victim), she can then launch a wide range of attacks such as persuading victims to click on malicious links, blackmailing, identity theft, and phishing. Some of the features provided by online social networks can be abused by attackers with the aim of launching automated reverse social engineering attacks. This form of attack can be categorized into three sub-groups, namely, recommendation-based, visitor tracking-based, and demographics-based reverse social engineering.

In the recommendation attack, the aim is to exploit the friend recommendations made by the social network to promote the fake profile of a fictitious user to the victim. The hope, from the attacker's point of view, is that the victim will be intrigued by the recommendation, and will attempt to contact the bogus profile that is under the attacker's control. In the visitor tracking attack, the aim is to trigger the target's curiosity by simply browsing her profile page. The notification that the page has been visited may be enough to attract the target to visit the attacker profile. Finally, in the demographic-based attack scenario, the attacker attempts to reach his victims by forging fake demographic or personal information with the aim of attracting the attention of users with similar preferences (e.g., similar musical tastes, similar interests, etc.).

These attacks highlight just a single facette of social networks. Other than friendship status and the involved level of trust, platform-based applications (Apps) represent another widely-used functionality with the potential to cause mischief. Probably everyone who has a Facebook profile has as least once been confronted with Farmville, Mafia Wars, birthday calendars or other apps through either news items on friends' walls or even direct requests by friends to use them. Although the times when third-party apps had unlimited access to a user's data are over by now, people still tend to willingly accept even boldest permission requests. One explanation for that behavior is that users often propagate their trust relationship to a friend directly to apps used by this friend [16]. Efforts to make users more aware of the privacy they are giving away might be a step into the right direction.

Another form of data exposure is presented by the possibility for third-party websites to interact with the social network by utilizing so-called plugins. Social plugins enable third-party websites to offer personalized content by leveraging

the social graph, and allow their visitors to seamlessly share, comment, and interact with their social circles [3]. For example, Facebook's Like button, probably the most widely deployed social plugin [1], enables users to leave positive feedback for the web page in which it has been embedded, share the page with their friends, and view their friends that have "liked" the page, along with the total number of "likes" from all visitors. Google's "+1" button [6] offers almost identical features to the Like button, while similar widgets are also available from other popular social networking sites such as Twitter and LinkedIn.

Social plugins have also been used for a wide variety of other applications including authentication. For example, instead of a web site implementing its own authentication system with user names and passwords, it may use a *social login* plugin offered by a social networking platform such as Facebook. In theory, this approach to authentication not only saves visitors from the burden of remembering one more password, but also gives them the opportunity to experience a personalized service from the web site based on their preferences and social circle.

Unfortunately, both technologies also bear an enormous risk to badly influence a user's privacy. In most cases, a visit to the target site is enough to identify the visitor, regardless of the actual interaction done with the plugin. Social login, on the other hand, enables third-party websites to access private information in a user's profile. A privacy leak not always anticipated by the user.

4 Outlook

In general, the evolution from traditional attacks to more specific forms that leverage social network information was logical. Where technological quirks, weaknesses and vulnerabilities acted as an enabler for traditional attack scenarios, relationships, trust and private information play an equally important role in social networks. Still, large-scale attacks with severe impact to the majority of participants of a social network have not been reported yet. In our opinion, the reason for this is twofold.

First of all, a social network is a strongly supervised and encapsulated structure where permissions are needed to carry out most actions (e.g. sending messages or posting comments). Misbehavior is promptly reported and the corresponding account blocked. Secondly, an attack, once implemented, does not necessarily yield the same results over time. In contrast to a deterministic, technological tool like a botnet or malware in general, the target in social networks are humans. And that bears the advantage of a certain capability to adapt to the circumstances. In the long run, even the most gullible user will be able to tell the difference between a legitimate friend request and a bogus one.

The greatest danger the users and participants of social networks have to face today, are privacy leaks. When the platforms have been introduced at first, they were designed as relatively closed environments which undoubtedly came with their own set of problems. In recent years, however, the progressive integration

of social networks into other branches made it increasingly difficult to track where personal information is used or where it can be accessed [8]. Even the tiny like-button discussed before, comes with its privacy issues, not to mention more advanced technologies like social authentication and other plugins.

For targeted attacks like spear phishing or social engineering, a social network is the perfect background. Even though the user is ultimately responsible for the amount of detail offered by her own presentation, researchers are prompted to raise the bar an attacker has to cross before successfully launching an attack. Previous research has proven the feasibility of keeping up or even staying ahead in the arms race. With ongoing effort it can be assured that it also holds true in the future.

References

[1] BuiltWith - Widgets Distribution. `http://trends.builtwith.com/widgets`.

[2] Facebook CSRF and XSS vulnerabilities. `http://www.john-jean.com/blog/advisories/facebook-csrf-and-xss-vulnerabilities-destructive-worms-on-a-social-network-350`.

[3] Facebook Plugins. `http://developers.facebook.com/docs/plugins/`.

[4] Facebook Security Phishing Attack In The Wild. `http://www.securelist.com/en/blog/208193325/Facebook_Security_Phishing_Attack_In_The_Wild`.

[5] Facebook Stats. `http://www.facebook.com/press/info.php?statistics`.

[6] Google +1 button. `http://www.google.com/+1/button/`.

[7] Sophos Facebook ID Probe. `http://www.sophos.com/pressoffice/news/articles/2007/08/facebook.html`, 2008.

[8] M. Balduzzi, C. Platzer, T. Holz, E. Kirda, D. Balzarotti, and C. Kruegel. Abusing Social Networks for Automated User Profiling. In *Recent Advances in Intrusion Detection*, pages 422–441. Springer, 2010.

[9] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All Your Contacts Are Belong to Us: Automated Identity Theft Attacks on Social Networks. In *18th International Conference on World Wide Web (WWW)*, 2009.

[10] R. Dhamija, J. D. Tygar, and M. Hearst. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, CHI '06, pages 581–590, New York, NY, USA, 2006. ACM.

[11] http://eu.battle.net/. Battle.net. 2 2012.

[12] https://steamcommunity.com/. The steam gaming community. 2 2012.

[13] http://www.origin.com/. Origin. 2 2012.

[14] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer. Social phishing. *Commun. ACM*, 50(10):94–100, 2007.

[15] G. Stringhini, C. Kruegel, and G. Vigna. Detecting spammers on social networks. In *ACSAC*, 2010.

[16] N. Wang, H. Xu, and J. Grossklags. Third-party apps on facebook: privacy and the illusion of control. In *Proceedings of the 5th ACM Symposium on Computer Human Interaction for Management of Information Technology*, page 4. ACM, 2011.