Social Network Security: A Brief Overview of Risks and Solutions

Edward Wang, ekw1@cec.wustl.edu (A project report written under the guidance of Prof. Raj Jain)

Download 

Abstract

In this study, we present the various aspects of social, network and physical security related with the use of social networks, by introducing the mechanisms behind each and summarizing relevant security studies and events related to each topic. It has been long understood that the widespread use of social networking sites can provide attackers with new and devastating attack vectors. In this study we attempt to dive deeper into each mode of security threat, as well as confirm the security risk associated with each topic by providing real world financial / social consequences. We recognize that while organizations and individuals may have legitimate business / personal uses for social networks, we recommend specific actions be taken to bolster stronger user awareness, more secure software designs as well as better organizational accountability.

Keywords

Social network security, social engineering, XSS, CSRF, DoS, stalking, OpenID, Facebook, twitter, LinkedIn, phishing, information theft, identity, identity hijacking, malware, worms, firewall, corporate security

Contents

Introduction

For a newcomer to the internet arena, social networking sites are an ever more popular way for people to stay connected. Some might even venture to say business opportunities are formed and lost online, as our web presence becomes an integral part of our personal lives. In an era where our online identity not only overshadows our actual identity, but other key financial and personal systems as well, the potential security risks associated with these social networks cannot be stressed enough.

Over the years, researchers and hackers alike have identified a handful of security risks ranging from people, process to application. The purpose of this study is to give a sweeping overview of the major security topics surrounding social networks today, and introduce the underlying mechanisms behind each. We follow up with some tangible consequences that each risk might have, and finally provide a direction to look at in terms of solutions.

Social Engineering

Information Leakage & Theft

Mechanism

Scope of Visibility

Most people when asked will agree that not everyone they know is their best friend; there are the mere acquaintances all the way to those with whom we share our deepest secrets, along with many shades in between. However the widespread phenomena of social networking sites has added new meaning to friends: two people are often "friends or not" (D, 2004). While social networks may not necessarily increase strong ties, it certainly does very little for weak ties. One may have a couple of close friends and thousands of distant friends, and a social network may simply categorize them all as "friends."

More contacts aren't necessarily a bad thing; the problem is who has access to our information? Social networking sites provide a certain level of access control, but most people do not take the effort to configure

these properly. As a result, everyone ends up with equal access rights. To make matters worse, oftentimes information travels through several hops of "friends," and by the idea of six degrees of separation it seems unreasonable to assume we are far from the bad guys.

Use of Real Names and Personal Information

As an added bonus, social networking sites contain information that is either mostly real or easily identified as fake. For the sole purpose of keeping up with friends in a seemingly trustworthy domain, people have very little incentive to falsify information on Facebook. The same idea goes for sites like MySpace and LinkedIn. See figure 1 for a recent study at Carnegie Melon University (Gross & Acquisti, 2005).

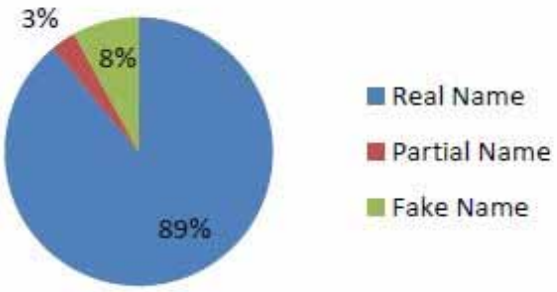

Figure 1: Percentage of CMU Facebook profiles and their respective name authenticity

Similar results exist for other sensitive information, such as birthdates, education history and hometown. In fact, a group of Taiwanese researchers have gone on to propose automated identification systems for name, age and education record inference on a different social network with good results (Lam, Chen, & Chen, 2008)

Breadth of Available Information

In the same CMU study, Gross and Acquisti go on to show the sheer amount of information available simply within the CMU Facebook realm. (Gross & Acquisti, 2005) Again, most users make very little effort to subdivide access privileges to different parts of their profile. By the same line of logic as names and birthdates, we have very little reason to doubt the validity of this information.

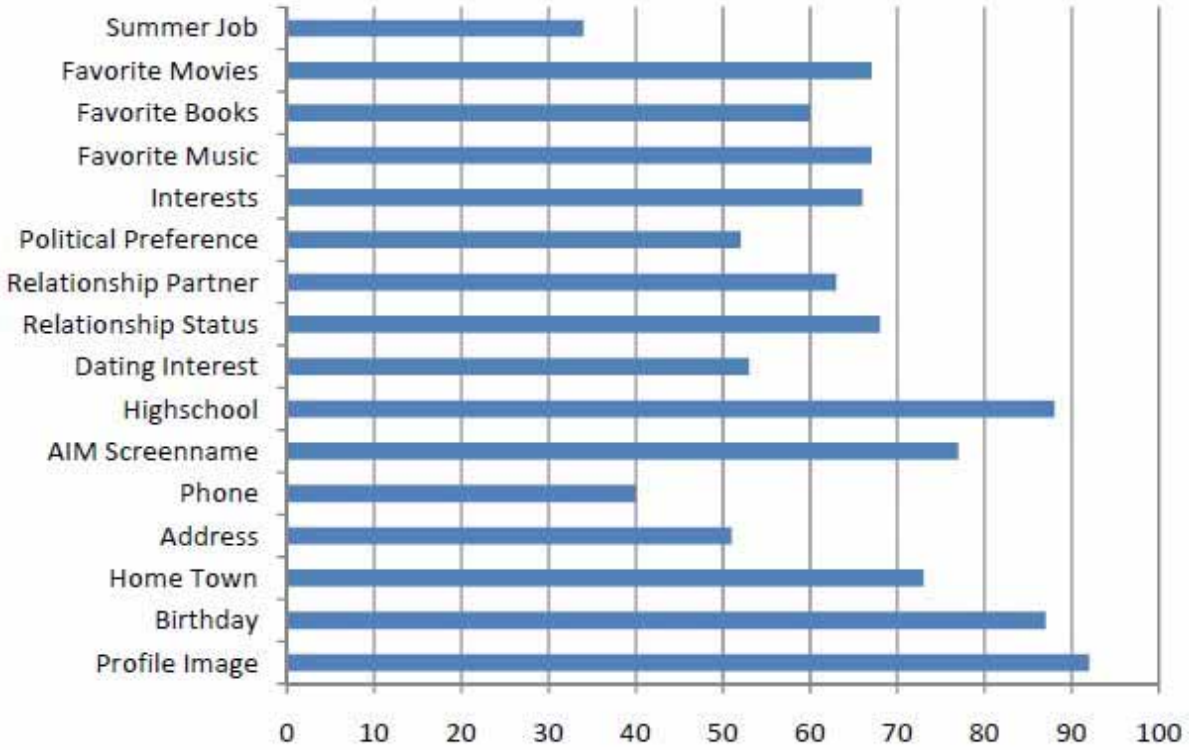

Figure 2: Percentages of CMU profiles revealing various types of personal information

Promiscuous Trust Relationships

Without taking into account the registration requirement relaxation for Facebook in the last year, users used to need a valid academic e-mail address in order to enroll. The bulk of Facebook users still operate under this assumption (as far as I can tell, the open Facebook doesn't bother the newcomers either), and we automatically trust whoever is in our network. Most campus networks are open and gaining a mail address is not difficult. Moreover, many users will gladly accept friend requests from people that aren't even in their network. (Jump, 2005)

As soon as a stranger connects with someone in a new network, he more or less inherits his friend's credentials when it comes to dealing with others in the same realm, giving him easy access to other users. As a bonus, this may allow a malicious user to circumvent realm-based privacy settings.

Data Protection Circumvention

Those who are more privacy conscious may choose to selectively disclose their information. However, as presented by figure 3, a malicious user is perfectly capable of utilizing Facebook's advanced search function and cross reference user IDs to gain access to hidden portions of a user's profile. (Gross & Acquisti, 2005)

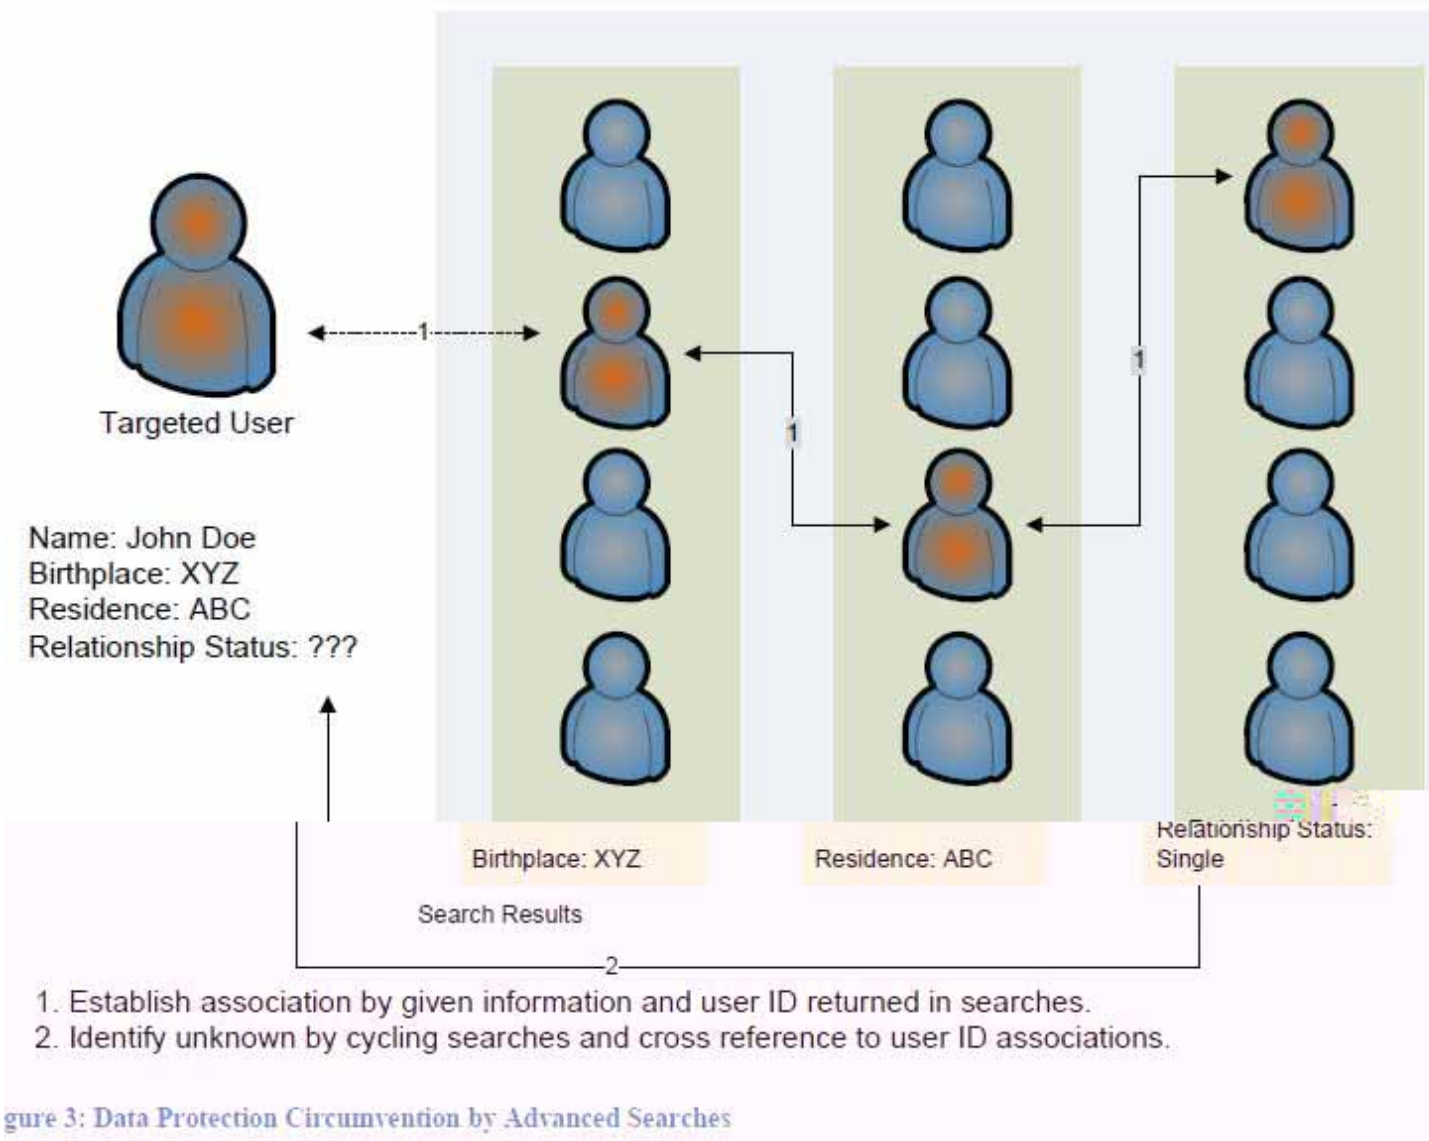

gure 3: Data Protection Circumvention by Advanced Searches

Re-Identification

It has been shown that most of the US population can be identified through the combination of ZIP code, birth date and gender (Sweeney). A quick glimpse over my own Facebook friends shows that most users provide this information.

Consequences

Information theft provides the jumping ground for a malicious user to mount a more targeted attack. Once the hacker has this information, he is free to engage in phishing, identity hijacking and many other forms of social attack. The information itself is also highly valuable in the eyes of advertisers and merchants, and could lead to other forms of privacy invasion and disturbances.

Possible Remedy

Lam, Chen and Chen provide some insightful recommendations to limit the scope of information propagation on social networks. They are abridged and reproduced in the following with my additional recommendations

and comments:

Personal Privacy Settings

Not only should users have greater control in assigning viewing privileges of personal information, the methodology should be more user friendly and streamlined. Although Facebook provides customizable groups and group-based authority control, setting one up weans away most of the users. A standardized model that group newly added friends to different pre-made privilege buckets would greatly enhance this process.

Browsing Scope Settings

The ability of users to view information across vast spans within a group should be limited. For example, detailed information should only be viewable up to a number of degrees away. LinkedIn provides good control in this category as a default, requiring authorization for anyone that is not a direct connection. Although this does not prevent anyone from hijacking a profile and gaining access through other means of social engineering (discussed later), it does provide a good starting ground to prevent widespread automated information harvesting.

Owner's Confirmation

Since much information is leaked not through the original profile but from innocent bystander comments, users should be given the authority to censor or control comments.

Phishing

Mechanism

As the name may suggest, phishing is the act of stealing sensitive information with a medium that has been designed to masquerade a legitimate source. (Merrian-Webster, 2009) For example, a malicious user could pretend to be a bank and send out mass emails asking for login credential verification. The user would most likely be redirected to a website that appears similar to the actual bank website that the hacker has constructed and be prompted to enter their credentials. Sophisticated phishes may leverage other techniques within social engineering and browser/server defect (such as XSS) to make the scheme look more plausible.

There are two key factors to a successful phishing scheme: 1) appearance of legitimacy and 2) trustworthiness of the delivery medium. After its humble beginning as the AOL IM user verification scheme (Stutz, 1998) and the ever so popular Nigerian scam (Wikipedia, 2009), modern phishing schemes rely on detailed personal information for targeted attacks. Studies show that spear phishing, or phishing with targeted personal information, can achieve up to an 80% success rate. (Bank, 2005) For a while we seemed to have this under control, with application level heuristic filters and network / firewall level domain screening. Enter social network sites such as Facebook (Facebook), MySpace (MySpace) and LinkedIn (LinkedIn). As previously stated, never before has the online population so willingly disclosed their personal information, such as phone numbers, addresses, interests and education history, that what used to be the end goal of major corporate espionage operations can now be collected at the click of a few buttons. Due to their business value and widespread acceptance, these sites are impervious to corporate firewalls. Combined with poor user education, powerful APIs and shoddy privacy settings, social network sites provide the perfect place to collect information and breed attack vectors.

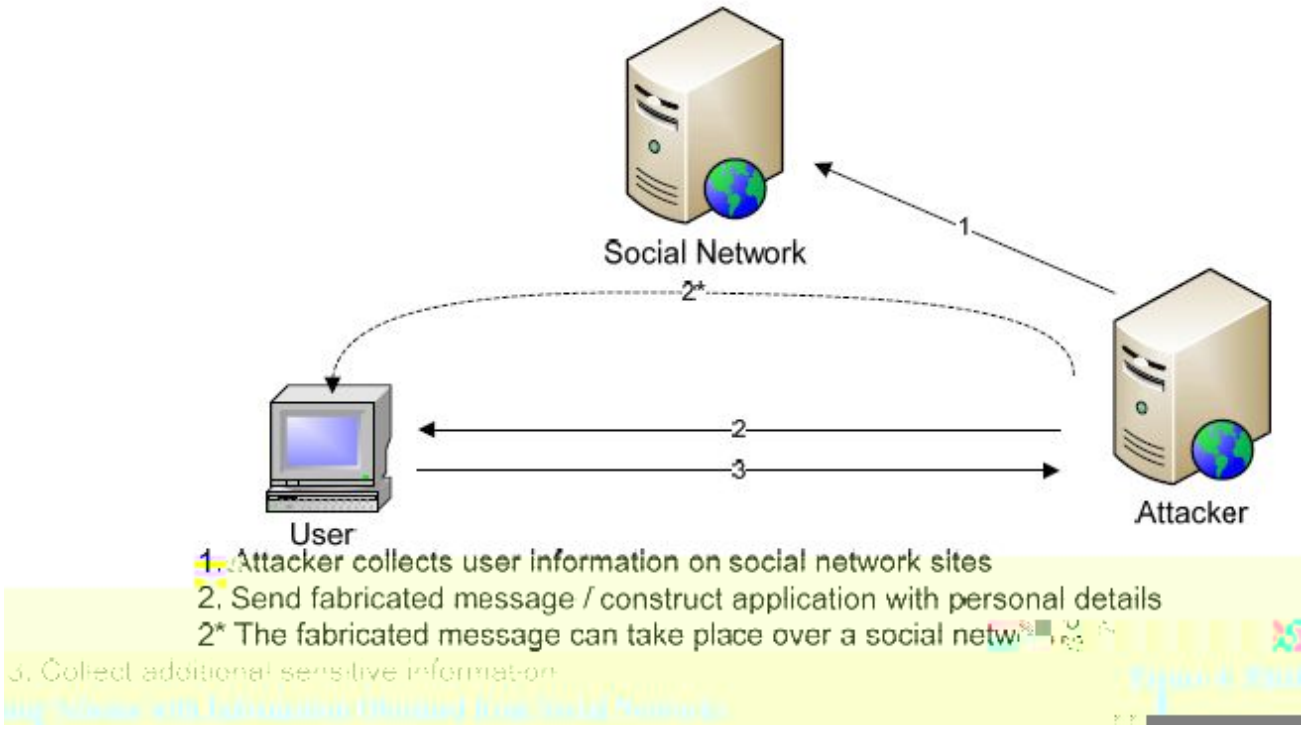

Typically, attackers leverage the available information on social networks to create highly specific and personalized messages. These messages are then either sent through the network directly or used in network applications to target their victims. As demonstrated by the 2006 MySpace attack , (websense, 2006) the high volume nature of social network and the built-in user trust for each other allowed a XSS worm to propagate without scrutiny.

Consequences

Although we have yet to see damages caused by phishing using social networks, the danger is real. The Facebook login credential scheme (Arrington, 2008) is a good example where, till this day, we still don't know the real motive of the attacker. We can only assume as more users jump on the social networking bandwagon (Facebook boasts 15,000 signups a day), the number and the severity of incidents should increase.

Possible Remedy

Users should be educated about the possible consequences phishing brings. Not only should users be sensitive about suspicious URLs and messages, they should follow strict policies regarding phishing site warning messages. Most importantly, users should review privacy settings in their social network applications and limit their exposure. At the application level, one may consider enforcing Microsoft Internet Explorer or Mozilla Firefox's phishing control. As an added protection, services like PhishTank (OpenDNS, 2009) and BlueCoat (Blue Coat Systems, 2009) provide network and application level blacklisting.

Identify Hijacking

Mechanism

A quick search of Google with the string "social network identity theft" returns more than enough results for a leisurely read. Cases ranging from Bryan Rutberg's identity theft on Facebook (Sullivan, 2009) to Lori Drew's famous yet tragic MySpace case (MarketingVOX, 2008) provide ample evidence that identity theft over social network is a big issue.

The idea of identity theft is perhaps nothing new, but what got tossed in the mix is most likely what was mentioned over and over again about the nature of social networks: the broad scope, the speed at which information travel, the ease for someone to establish an identity, and finally the disproportionate amount of trust people place on others on the network.

First, as demonstrated in Bryan Rutberg's case on Facebook and Nathan / Shawn's ability to masquerade as Marcus Ranum in Black Hat / Def Con (Hamiel & Moyer, 2007), it takes mere minutes and hours, not days and months, for a successful attack to complete. The attackers were able to reach all of Bryan's friends within seconds, just as Nathan / Shawn were able to seek out all potential leads and get approval within hours. The victim may not realize the act until long after the attack is complete, or may not have enough time to respond due to the speed disconnect between formalized response and the online attack.

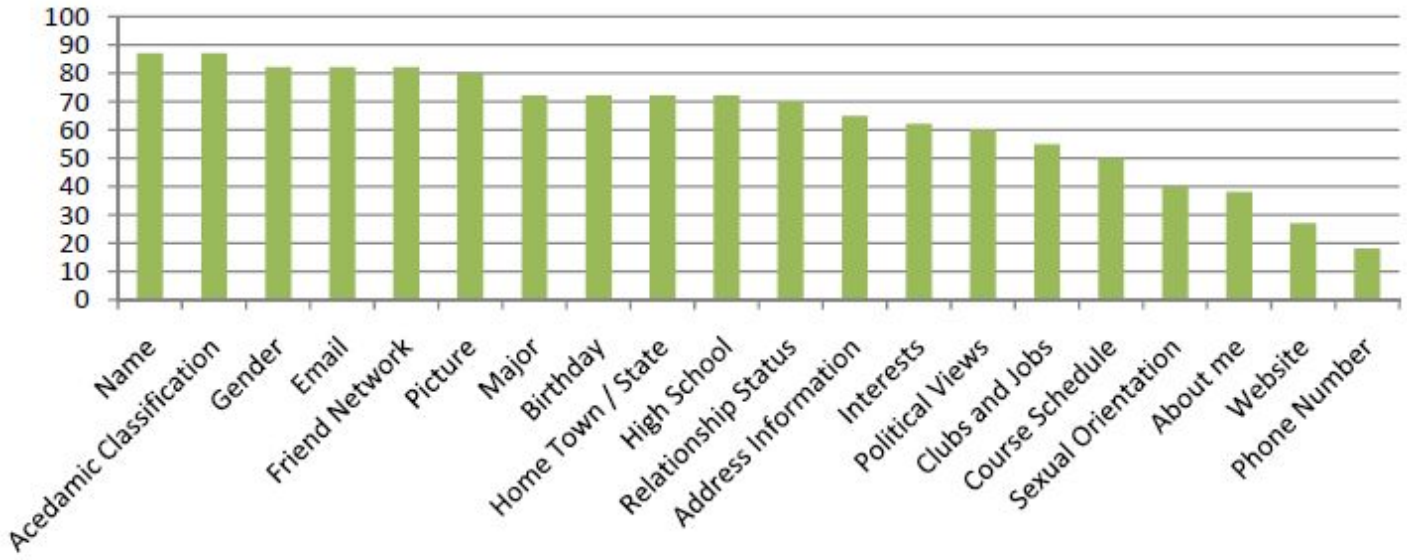

Figure 5: Primary Identity Information Disclosure on Facebook

Second, due to poor identity authentication and the availability of personal information on social networks, attackers have plenty of opportunity to carry out a targeted attack. It only took Nathan and Shawn few hours to track down the necessary information to build a convincing profile of Marcus Ranum on LinkedIn. He doesn't even have any previous social networking profile! Consider another study done to explore the primary identity information disclosure on Facebook as shown in figure 5 (Stutzman, 2006); building a convincing profile of someone is very much a trivial task.

Finally, people trust whatever their friends on social networks say, as well as those who aren't even really friends. (Walther, Heide, Kim, Westerman, & Tong, 2008) A malicious user may, through hijacking existing accounts or by accounts created through information mining, convincingly execute plots on others with little detection.

Consequences

Identity theft cost the average victim $5,720 per incident on average in 2007, up from $740 in 2003. (Claburn, 2007) While a study conducted by Utica College's Center for Identify Management and Information Protection

(CIMIP) found the number to be much higher , the magnitude is out of the scope of this discussion. What is important to realize is that petty identity theft may cost is money at the best, it can escalate very quickly from there. So far we have only seen petty crimes and incidents. Considering the power shown by Nathan / Shawn, the amount of damage a professional organization can cause with social network identity theft is unmanageable.

Possible Remedy

Email Verification

Although Email security is largely dependent on how well an organization locks down the rest of the network, corporate social networks should still check for valid addresses as a first line of defense. Users should also be educated to check for valid addresses in profiles that they are interacting with.

CA and one-point login for users

The whole idea about the internet has been founded on anonymity, but this notion might not be the best with the amount of trust we place on each other over the web today. Anyone can sign up for accounts either in someone else's name and perform activities without the original user's knowledge. If instead everyone is assigned a certificate, very much like a driver's license or SSN in real life, that ties into our online credentials and positively identifies us to each other over the web, the possibility of identity theft would drastically reduce.

Another attack vectors for account hijacking comes from weak passwords, or password shared across many sites. Solutions like VeriSign's Personal Identity Portal (PIP) (VeriSign, 2009) and OpenID (OpenID, 2009) may solve this problem: PIP can be configured to perform two factor authentications and OpenID provides a central location for high grade security logon credentials. A malicious user would be forced to obtain not only the password but the key fob (often a cell phone) as well. A commercial entity like VeriSign can provide quick and deterministic actions in the event of an account theft, and OpenID would be a great way for users to cut off unauthorized access to all affected accounts.

User Education

The old saying holds: users should take anything on the internet with a grain of salt. Take precaution when dealing with people over social networking sites, even those that claim to be someone we know. Special care should be used to verify sensitive information received. Last but not the least, use safe passwords and don't use the same password on every site.

Physical Security

Stalking

Mechanism

With the granularity of the information available on social networking sites, it may be possible for a malicious user to figure out where someone is going to be at a certain time. Many students provide information such as residence and full class schedules on Facebook, combined with status updates stating what they are doing or where they are. For those adults not in schools, frequent Twitter updates (Higgins, The Seven Deadliest Social

Networking Hacks , 2008) combined with social network cross-referencing can account for the same effect.

Consequences

The ability for malicious users to figure out where a target is physically is very dangerous, not only on a privacy issue but more so on a personal safety level. It opens up opportunities for burglary, assault and kidnapping.

Possible Remedy

Higgins recommends toning down details posted on sites such as MySpace, Facebook and Twitter regarding one's whereabouts. Status updates should not contain specifics, and should only be posted after the fact.

Malware

Cross-Site Reference Forgery (CSRF) & Cross-Site Scripting (XSS)

Mechanism

What is CSRF

CSRF, quite simply, is a malicious user exploiting certain site's trust for an existing authenticated session with a victim's browser. While most web applications require credentials for the initial authentication, many subsequent actions are simply authenticated in background via limited-lifetime cookies. If an application does not require re-authentication and the cookie is still valid, a crafty attacker can trick the victim's browser to execute actions on a site's application with existing valid credentials. (Wikipedia, 2009)

There are many ways a CSRF attack vector can be delivered. As a common requirement, the attacker has to trick the victim into executing a POST or GET request. This can be done via an IMG SRC tag, a malicious link, or more complicated mechanisms that are out of the scope of this discussion.

CSRF Take Two: Logon CSRF

As web applications start to apply session and link specific tokens, one area overlooked is often the logon process. Since there are no active credentials to bind a session token, most sites do not check for valid GET/POST requests during logon. Facebook falls into this category; while subsequent pages all have unique IDs, the logon process only check for a valid REFER_ID. (Barth, Jackson, & Mitchell, 2008) Flaws like this allow an attacker to trick a victim to authenticate into web applications using the attacker's credentials, making the victim's information recoverable at a later date.

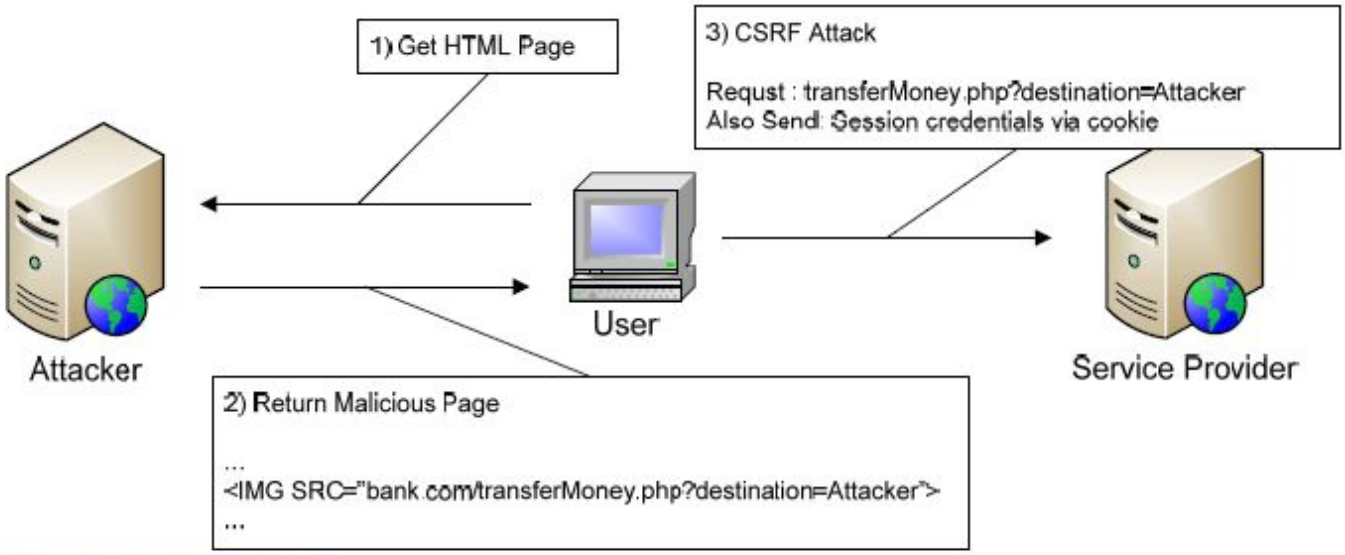

Figure 6: Typical CSRF Scheme

A variant of CSRF: DoS Opportunities

The avid reader might quickly realize that the power to control what thousands, if not millions of browsers are viewing is a powerful tool. In fact, a group of researchers distributed a simple picture of the day application on Facebook, and had over 300 views / hour within days. (Athanasopoulos, et al., 2008) A malicious user can easily insert multiple requests for resource heavy objects on a victim's server, and perform a powerful and effective DoS attack. *What is XSS*

While CSRF exploits the server's trust for a client's browser, XSS exploits a client's browser for the server. More strictly speaking, XSS takes advantage of the often-overlooked encoding rules that browsers display pages. An attacker may insert executable code into a posting or image tag. Vulnerable systems often do not discriminate between readable content and code, and will faithfully represent them to other clients. A victim's browser will then execute the code portion of the posting without user's knowledge. (Wikipedia, 2009)

The two most often executed schemes include JavaScript and browser state mutations. JavaScript allows for a range of attacks from replacing content as seen on a victim's screen to executing additional embedded code, while browser state mutations allow the attacker to gain access to a victim's cookies and other vital personal information.

CSRF & XSS + Social Networking Sites

While a powerful combo, the biggest problem hackers faced with these attacks was distribution. Not only do they need a large amount of exposures, they also needed a victim's trust to perform some initial actions (such as clicking a link.) Enter social networking sites as a rescue: 3rd party derived contents can be rapidly and reliably distributed to almost everyone, and people are more inclined to click on random links than ever before. Despite the filtering and API protections that developers are trying too hard to erect, the danger is still real. Billy Rios and Raghav Dube, senior security analysts at Ernst & Young's Advanced Security center demonstrated JavaScript code that can steal a user's credentials and attack another site over social networking sites. (Higgins, Killer Combo: XSS + CSRF, 2007)

Consequences

Besides the direct financial consequences from CSRF and XSS, these attacks are the springboard for malicious users to gather sensitive information about potential victims for more targeted attacks. Given the relative

newness and complexity of knowledge required we have yet to see sophisticated application of these attacks, however the potential is there and the harm should not be ignored.

Possible Remedy

Defenses against CSRF and XSS are ongoing research and deserve a far deeper analysis than what I will be able to get into here. However, the general ideas are as follows.

CSRF: Secret Token Validation

Each client-server request should be protected via unique tokens, preferably seeded by nonces. To reduce the possibility of DoS attacks, a token should be somehow bindable to a user's session without having the server keep excessive amount of state information. A simple hash of a username and session identifier should be avoided, as it allows an attacker to acquire a valid session identifier. (Barth, Jackson, & Mitchell, 2008)

Several software packages exist as plug-ins for server side defense against CSRF. NoForge (Technical University of Vienna, 2007) is an interesting example, as it modifies every single link as it is serialized onto the network and appends a unique token. Several issues exist with this approach that warrants a more fundamental solution. At the end of the day, the issue boils down to good integration of token management at the initial software design phase.

CSRF: Referrer Validation

While pre-validation conditions warrant the use of referrer checking, many conditions render is approach ineffective. Many browsers strip referrer tags by default when hopping between HTTP and HTTPS states, not to mention most corporate networks will do this to retain internal privacy. As a result, most sites will not enforce strict referrer checking. Instead, consider setting up pre-sessions prior to validation and transfer the session once validation occurs.

XSS: Strict Character Validation / Encoding Rules

As a rule thumb, special or escape characters should always be screened at input. Server outputs should be strictly encoded to ensure that text is displayed as text, and only authorized blocks are allowed to execute as code. There are complete implementation guides available regarding this topic that covers all the bases much better than what I can place here.

Conclusion and Advice

In this paper we have shown the various means a malicious user can leverage social networks as an attack vector. Ranging from information harvesting to sophisticated scripting, social networks pose a real yet elusive security threat. There is a vast amount of sensitive personal information available on social networks, and the lack of proper security levels at the user level make these sprawling applications an ideal sandbox for attackers. As shown by the Facebook account hijacking incident, attack can bring real financial damages. Other damages are also possible, only to go up in severity.

The best countermeasure still seemed to be user education – users need to know the consequences of publishing detailed personal information, they also need safer ways to handle cyberspace events. Whether it is friend requests or potential phishing messages, users should understand the proper ways to safeguard their own well being, as well as those around them in a group environment. Other methods such as better security

measures from social networking application and end-user software, barrier-level protection and greater corporate accountability can all work together to curb potential attacks.

It would appear that we are addicted to social network sites, and the momentum is only picking up. With the immense business and personal opportunities offered by these venues, it is hard for corporations and individuals to completely shut these systems out. The world of securing these popular and ever-changing applications is only starting, and it is exciting to see where the future will take us.

Bibliography

- Arrington, M. (2008). Phishing For Facebook. Retrieved from http://www.techcrunch.com/2008/01/02/phishing-for-facebook/
- Athanasopoulos, E., Makridakis, A., Antonatos, S., Antoniades, D., Ioannidis, S., Anagnostakis, K., et al. (2008). Antisocial Networks: Turning a Social Network into a Botnet. Lecture Nots in computer Science .
- Bank, d. (2005, August). 'Spear Phishing' Tests Educate People About Online Scams. Retrieved from http://online.wsj.com/public/article/SB112424042313615131-z_8jLB2WkfcVtgdAWf6LRh733sg_20060817.html?mod=blogs
- Barth, A., Jackson, C., & Mitchell, J. (2008). Robust Defenses for Cross-Site Request Forgery. 15th ACM Conference on Computer and Communications Security.
- Blue Coat Systems. (2009). Blue Coat. Retrieved from http://www.bluecoat.com/node/1323
- Claburn, T. (2007, October). Identity Theft: Costs More, Tech Less . Retrieved from http://www.informationweek.com/news/internet/showArticle.jhtml?articleID=202600312
- D, B. (2004). Friendster and publicly articulated social networking. Conference on Human Factors and Computing Systems.
- Facebook. (n.d.). Retrieved from facebook.com
- Gross, R., & Acquisti, A. (2005). Information Revelation and Privacy in Online Social Networks (The Facebook Case). ACM Workshop on Privacy in the Electronic Society .
- Hamiel, N., & Moyer, S. (2007). Satan is on My Friends List. Retrieved from http://www.hexsec.com/docs/Satan_Blackhat_Defcon.pdf
- Higgins, K. J. (2007, March). Killer Combo: XSS + CSRF. Retrieved from http://www.darkreading.com/security/app-security/showArticle.jhtml?articleID=208804442
- Higgins, K. J. (2008, August). The Seven Deadliest Social Networking Hacks . Retrieved from http://www.darkreading.com/security/app-security/showArticle.jhtml;jsessionid=GCLXYJXQFTCDEQSNDLPSKHSCJUNN2JVN?articleID=211201065&pgno=2&queryText=&isPrev
- Jump, K. (2005). A new kind of fame. The Columbian Missourian .
- Lam, I.-F., Chen, K.-T., & Chen, L.-J. (2008). Involuntary Information Leakage in Social Network Services. 3rd International Workshop on Security: Advances in Information and Computer Security.
- LinkedIn. (n.d.). Retrieved from linkedin.com
- MarketingVOX. (2008, November). Watershed Ruling in MySpace Suicide Case May Criminalize Fake 'Net Personas'. Retrieved from http://www.marketingvox.com/watershed-ruling-in-myspace-suicide-case-may-criminalize-fake-net-personas-042175/
- Merrian-Webster. (2009). Phishing. Retrieved from Merrian-Webster: http://www.merriam-webster.com/dictionary/phishing
- MySpace. (n.d.). Retrieved from myspace.com
- OpenDNS. (2009). PhishTank. Retrieved from phishtank.com
- OpenID. (2009). OpenID. Retrieved from http://openid.net/
- Stutz, M. (1998). AOL: A Cracker's Paradise? Retrieved from http://wired-vig.wired.com/science

/discoveries/news/1998/01/9932
- Stutzman, F. (2006). An Evaulation of Identity-Sharing Behavior in Social Network Communities.
- Sullivan, B. (2009, January). Facebook ID Theft Targets 'Friends'. Retrieved from http://redtape.msnbc.com/2009/01/post-1.html
- Sweeney, L. Uniqueness of simple demographics in the U.S. Population. Carnegie Mellon University, Laboratory for Internal Data Privacy.
- Technical University of Vienna. (2007). NoForge. Retrieved from http://www.seclab.tuwien.ac.at /projects/noforge/
- VeriSign. (2009). Personal Identity Portal. Retrieved from https://pip.verisignlabs.com/
- Walther, J. B., Heide, B. V., Kim, S.-Y., Westerman, D., & Tong, S. T. (2008). The Role of Friends' Appearance and Behavior on Evaluations of Individuals on Facebook: Are We Known by the Company We Keep? Human Communication Research .
- websense. (2006). MySpace XSS QuickTime Worm. Retrieved from http://securitylabs.websense.com /content/Alerts/1319.aspx
- Wikipedia. (2009). Advance-fee fraud. Retrieved from http://en.wikipedia.org /wiki/Nigerian_scam#cite_note-1
- Wikipedia. (2009). Cross-site request forgery. Retrieved from http://en.wikipedia.org/wiki/CSRF
- Wikipedia. (2009). Cross-site scripting. Retrieved from http://en.wikipedia.org/wiki/Cross-site_scripting

---

List of Acronyms

- XSS - Cross site scripting
- CSRF - Cross site reference forgery
- DoS - Denial of service
- CMU - Carnegie Mellon University
- ID - Identification
- US - United States
- ZIP - Zone improvement plan
- AOL - America Online
- IM - Instant message
- CIMIP - Utica College's Center for Identify Management and Information Protection
- PIP - Personal Identity Portal, VeriSign

Last Modified: April 19th, 2009
This and other papers on latest advances in network security are available on line at http://www.cse.wustl.edu /~jain/cse571-09/index.html


Back to Raj Jain's Home Page