

The Devil is Phishing: Rethinking Web Single Sign-On Systems Security

(Extended Abstract)

Chuan Yue (cyue@uccs.edu), *University of Colorado Colorado Springs, USA*

Abstract

One significant trend in online user authentication is using Web Single Sign-On (SSO) systems. Especially, open Web SSO standards such as OpenID and OAuth are rapidly gaining adoption on the Web, and they enable over one billion user accounts. However, the large-scale threat from phishing attacks to real-world Web SSO systems has been significantly underestimated and insufficiently analyzed. In this paper, we (1) pinpoint what are really unique in Web SSO phishing, (2) provide one example to illustrate how the identity providers (IdPs) of Web SSO systems can be spoofed with ease and precision, (3) present a preliminary user study to demonstrate the high effectiveness (20 out of 28, or 71% of participants became “victims”) of Web SSO phishing attacks, and (4) call for a collective effort to effectively defend against the insidious Web SSO phishing attacks.

Keywords: Web Single Sign-On (SSO), Security, User Authentication, Phishing, Password

1 Introduction

Web Single Sign-On (SSO) systems allow users to sign in multiple relying party (RP) websites using one single identity provider (IdP) account; therefore, users are relieved from the huge burden of registering many online accounts and remembering many passwords. For example, a user can sign in many RP websites such as foxnews.com and sears.com directly using an IdP account such as a Gmail or a Facebook account.

This usability advantage along with the deployment incentives such as user profile sharing provided by IdPs make a major contribution to the rapid adoption of open Web SSO standards such as OpenID and OAuth in recent years. There are one billion OpenID enabled user accounts and over 50,000 RP websites [21], and a few leading IT companies including Google, Facebook, Microsoft, and Yahoo are OpenID IdPs. OAuth 2.0 autho-

rization framework has also been widely supported by IdPs and adopted by a large number of RP websites [15].

Researchers have performed formal security analysis of some Web SSO protocols [1, 2], and have identified logic flaws and implementation vulnerabilities of many deployed Web SSO systems [9, 11]. OpenID and OAuth designers and researchers have considered phishing attacks before [10, 14, 17, 18]. However, those considerations did not provide in-depth insights on the uniqueness of Web SSO phishing, and thus have not been taken seriously by the community in the process of promoting the Web SSO standardization and adoption [16, 20].

In this paper, we argue that the large-scale threat from phishing attacks to real-world Web SSO systems has been significantly underestimated and insufficiently analyzed. In particular, we support this argument in Section 2 by (1) pinpointing what are really unique in Web SSO phishing, (2) providing one example to illustrate how the IdPs of Web SSO systems can be spoofed with ease and precision, and (3) presenting a preliminary user study to demonstrate the high effectiveness (20 out of 28, or 71% of participants became “victims”) of Web SSO phishing attacks. Finally, in Section 3, we call for a collective effort to effectively defend against the insidious Web SSO phishing attacks.

2 Web SSO Phishing Attacks

We now support our argument that the large-scale threat from phishing to real-world Web SSO systems has been significantly underestimated and insufficiently analyzed.

2.1 What are Unique in Web SSO Phishing

In comparison with traditional phishing [8], Web SSO phishing possesses at least three unique characteristics: (1) the value of IdP accounts is highly concentrated, (2) the attack surface area is highly enlarged, and (3) the difficulty of phishing detection (either by algorithms or by

users) is highly increased. All these characteristics make Web SSO phishing more profitable and insidious than traditional phishing, and make it very challenging to defend against the attacks. Therefore, phishers have greater incentives to focus on attacking users' IdP accounts.

2.1.1 Highly Concentrated Value of IdP Accounts

The majority of the RP websites use the authentication or authorization services provided a few major IdPs such as Google, Facebook, Microsoft, and Yahoo. In general, a user's accounts such as Google and Facebook accounts are already highly valuable by themselves. Using these accounts as IdP accounts has further concentrated their value because a compromised IdP account also allows attackers to impersonate the victim on a large number of RP websites. As a result, phishing IdP accounts becomes much more profitable than before.

2.1.2 Highly Enlarged Attack Surface Area

Traditionally, phishers mainly use spoofed emails to lure users to the phishing websites. A spoofed email provides the *first-level context* (e.g., asking for account verification or update) to entice users to click a phishing URL, and the spoofed phishing website further provides the *second-level context* (e.g., with look and feel similar to a targeted real website) to lure users to submit their login credentials. In terms of this first-level context, the success of traditional phishing is limited by two main constraints. One is that if phishing emails are suspicious, the majority of users would not click phishing URLs and visit the phishing websites [3, 5]. The other is that a large number of phishing emails are captured by spam filters [12], thus cannot even reach users in the first place.

Now with Web SSO, clicking a button that represents an IdP (e.g., Google or Facebook) to visit the IdP's login webpage becomes a common practice. Thus, phishers are freed from these two main constraints – they can host their own “legitimate” RP websites or webpages such as for shopping or gaming (note phishers can still spoof legitimate RP websites) and lure users to visit by posting URLs everywhere (e.g., Web forums, blogs, and advertisements). In other words, the first-level context can be provided by whatever means in addition to spoofed emails. Therefore, the attack surface area is highly enlarged in Web SSO phishing, heightening the chance of visiting phishers' websites or webpages by many users.

2.1.3 Highly Increased Phishing Detection Difficulty

On their RP websites or webpages, phishers only need to spoof the login webpages of IdPs, and will only display a spoofed IdP login webpage if a user-initiated click event occurs on the corresponding button. Figure 1 illustrates

a typical example of Web SSO login buttons hosted on many RP websites. Similar buttons can also be hosted on the RP websites owned or controlled by phishers.



Figure 1: A typical example of Web SSO login buttons.

Figure 2(a) illustrates a typical real Google login webpage displayed after a user clicks the Google button on many RP websites such as foxnews.com and sears.com. Currently, all such IdP login webpages are displayed in a popup window created by the JavaScript `window.open()` method, which is supported in all major browsers.

Such a *click-and-popup* user interaction style is the root cause of the highly increased phishing detection difficulty. A popup IdP login window, although displaying the EV-SSL (Extended Validation SSL) icon and an HTTPs URL address, can be spoofed with ease and precision as will be exemplified soon in Section 2.2. A spoofed popup login window does not need to correspond to a real URL address, thus can make the results of a large number of URL-heuristics-based automatic phishing detection algorithms (e.g., [4, 6, 7, 12, 13]) either inaccurate or incorrect. Meanwhile, as will be shown in Section 2.3, the look and feel of a spoofed popup login window can also deceive many users.

2.2 Spoofing IdPs with Ease and Precision

Figures 2(b) and 2(c) illustrate the examples of the spoofed login webpages for Google and Facebook, respectively. Each webpage is popped up when a user clicks the corresponding button shown in Figure 1. We created these webpages mainly using HTML, CSS (Cascading Style Sheets), and JavaScript. Spoofed login webpages of other IdPs can also be created.

The essential trick is that such a spoofed login webpage is not contained in a real popup browser window – it is indeed contained in an HTML `<div>` (i.e., division) element, which is supported in all major browsers. This `<div>` element has a larger CSS “*z-index*” value than its parent element, thus making it rendered on the top of the current webpage like in a real popup window. We used the dialog widget in jQuery library (<http://jquery.com/>) to create this type of fake “popup” window.

Another important trick is on spoofing the EV-SSL icon and the HTTPs URL address in the `<div>` element. This can be done by copying a complete snapshot of the icon and the URL address (as in Figures 2(b) and 2(c)) from a real IdP login window such as Figure 2(a). Alter-



Figure 2: (a) A typical real Google login page, (b) A spoofed Google login page, (c) A spoofed Facebook login page.

natively, the EV-SSL icon can be an image, but the URL address can be a string in a non-editable text input element. Either way, we use CSS to adjust the position and size of the spoofed icon and URL, making them look like real ones. We can further associate a click event handler to the icon for displaying copied EV certificate information, thus further making them feel like real ones. Since the spoofed login webpages do not appear in real popup windows and do not really have URL addresses, it is very difficult for URL-heuristics-based automatic phishing detection algorithms to properly detect the scams.

The browser name in the title of a spoofed login webpage (“Google Chrome” as in Figures 2(b) and 2(c)) is detected and displayed by JavaScript for each individual user. For the other elements in the two spoofed login webpages, we mainly copied and adopted the HTML and CSS contents from the real Google and Facebook login popup windows. The only key difference is the *action* attribute of the login form – on the spoofed webpages, the submitted login credentials will be sent to phishers.

If we carefully compare the real and spoofed Google login webpages in Figure 2(a) and Figure 2(b), we can still observe differences such as color schemes and title bars. All these differences can be further addressed using the HTML, CSS, and JavaScript techniques that we used for the other parts of the webpages.

2.3 A Preliminary User Study

To measure whether regular users can properly detect the spoofed IdP login webpages, we conducted a user study in early May 2013. The study was pre-approved by the IRB (Institutional Review Board) of our university.

2.3.1 Participants and Procedure

Twenty-eight (which offers a reasonably tight confidence interval [19]) adults, 14 females and 14 males, participated in our user study. They were voluntary stu-

dents and faculty/staff members randomly recruited in our campus library and bookstore, and they came from 19 departments of our university. Twenty-six participants were between ages of 18 and 30, and two participants were over 30 years old. We did not screen participants based on their Web browsing experiences. We did not provide monetary compensation to the participants.

We created a simple shopping website with the SSO login buttons shown in Figure 1 and with the corresponding spoofed Google and Facebook login webpages shown in Figures 2(b) and 2(c). We asked each participant to log into the shopping website using either a Google or a Facebook test account provided by us, and using one of the three most popular browsers (i.e., Google Chrome, Firefox, and IE) with the latest versions.

2.3.2 Results and Analysis

We collected results through two questionnaires (pre-procedure and post-procedure) and observation. From the results of the pre-procedure questionnaire, we found that 27 participants use browsers daily and one participant uses browsers weekly. In another question, by looking at the Web SSO example webpage of one RP website (sears.com), 22 participants answered that they have Web SSO experience and have logged into some other websites using their Gmail or Facebook account before.

After performing the Web SSO procedure on our shopping website, each participant answered two main questions one by one in the post-procedure questionnaire: **Q1**: “Is that Gmail or Facebook login page a genuine one?”, and **Q2**: “Have you heard about phishing attacks?”. Figure 3 illustrates the Venn diagram of answers to Q1 and Q2. We can see that 20 (or 71% of) participants answered “Yes” to Q1, and 24 (or 86% of) participants answered “Yes” to Q2. Furthermore, 17 (or 61% of) participants who have heard about phishing attacks were deceived by the spoofed login webpages. These results indicate that the spoofing techniques presented in

Section 2.2 are very effective, and the success rate of Web SSO phishing can be much higher than that of traditional phishing (around 10% as reported in [4, 5]).

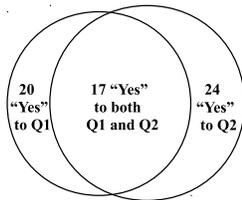


Figure 3: The Venn diagram of answers to Q1 and Q2.

Many of the 20 participants who became “victims” (i.e., answered “Yes” to Q1) commented that our spoofed login webpages look real and credible, and/or they saw similar login webpages before. Among the 8 participants who answered “No” to Q1, one explained that asking for Google/Facebook information is suspicious, and this participant indeed did not use Web SSO before; another explained that the shopping website is too simple and suspicious; the third explained that the test Gmail account is suspicious; the other 5 participants mainly explained that the color and menu bar on the login webpages are suspicious. As explained in Section 2.2, all these differences can be further addressed.

We observed that 16, 6, and 6 participants used Google Chrome, Firefox, and IE to perform the Web SSO procedure, respectively. Meanwhile, 23 and 5 participants used our test accounts to interact with the spoofed Google and Facebook login webpages, respectively. The most astonishing observation is that *none* of the 28 participants clicked the spoofed EV-SSL icon and the HTTPS URL address, indicating that users rely more on *look* than *feel* to identify the credibility of websites.

We explained in detail our Web SSO phishing attack to all the 28 participants just before they left the study. All of them appreciated our explanation, and the 20 “victims” were also very surprised by the fact that those two IdP login webpages are not real ones.

3 Discussion and Conclusion

To effectively defend against the insidious Web SSO phishing attacks, we call for a collective effort from browser vendors, IdPs, RPs, and users. We discussed in Section 2.1.3 that the *click-and-popup* user interaction style is the root cause of the highly increased phishing detection difficulty for algorithms and users. First, we believe this root cause must be addressed. We are exploring some smooth Web content and context switching techniques along this direction.

Second, IdPs should further improve their anti-phishing capabilities. For example, they may need to actively adopt a two-factor authentication approach (e.g., Google’s 2-step verification) and incentivize users to use

it. Note that two-factor authentication mitigates the consequence of Web SSO phishing, but does not prevent or detect it; therefore, the IdP password factor is still at risk.

Last but not least, our user study results clearly demonstrate that users should be educated and trained to understand and identify Web SSO phishing. We thank anonymous reviewers for their insightful comments. We sincerely welcome further suggestions and discussion.

References

- [1] ARMANDO, A., CARBONE, R., COMPAGNA, L., CUELLAR, J., AND TOBARRA, L. Formal analysis of saml 2.0 web browser single sign-on: breaking the saml-based single sign-on for google apps. In *ACM FMSE* (2008).
- [2] BANSAL, C., BHARGAVAN, K., AND MAFFEIS, S. Discovering concrete attacks on website authorization by formal analysis. In *IEEE CSF* (2012).
- [3] DOWNS, J. S., HOLBROOK, M. B., AND CRANOR, L. F. Decision strategies and susceptibility to phishing. In *SOUPS* (2006).
- [4] GARERA, S., PROVOS, N., CHEW, M., AND RUBIN, A. D. A framework for detection and measurement of phishing attacks. In *WORM* (2007).
- [5] JAKOBSSON, M., AND RATKIEWICZ, J. Designing ethical phishing experiments: a study of (ROT13) rOnl query features. In *WWW* (2006).
- [6] LUDL, C., MCALLISTER, S., KIRDA, E., AND KRUEGEL, C. On the effectiveness of techniques to detect phishing sites. In *DIMVA* (2007).
- [7] MA, J., SAUL, L. K., SAVAGE, S., AND VOELKER, G. M. Beyond blacklists: learning to detect malicious web sites from suspicious urls. In *SIGKDD* (2009).
- [8] RACHNA DHAMIJA, J.D. TYGAR, AND MARTI HEARST. Why phishing works. In *CHI* (2006).
- [9] SUN, S.-T., AND BEZNOSOV, K. The devil is in the (implementation) details: an empirical analysis of oauth sso systems. In *CCS* (2012).
- [10] SUN, S.-T., POSPISIL, E., MUSLUKHOV, I., DINDAR, N., HAWKEY, K., AND BEZNOSOV, K. What makes users refuse web single sign-on?: an empirical investigation of openid. In *SOUPS* (2011).
- [11] WANG, R., CHEN, S., AND WANG, X. Signing me onto your accounts through facebook and google: A traffic-guided security study of commercially deployed single-sign-on web services. In *the IEEE Symposium on Security and Privacy* (2012).
- [12] WHITTAKER, C., RYNER, B., AND NAZIF, M. Large-scale automatic classification of phishing pages. In *NDSS* (2010).
- [13] ZHANG, Y., HONG, J., AND CRANOR, L. CANTINA: A content-based approach to detecting phishing web sites. In *WWW* (2007).
- [14] OAuth 2.0 Threat Model and Security Considerations. <http://tools.ietf.org/html/rfc6819>.
- [15] OAuth Introduction. <http://oauth.net/about/>.
- [16] OpenID Authentication 2.0 - Final. http://openid.net/specs/openid-authentication-2_0.html.
- [17] OpenID Phishing Brainstorm. http://wiki.openid.net/w/page/12995216/OpenID_Phishing_Brainstorm.
- [18] OpenID: Phishing Heaven. <http://www.links.org/?p=187>.
- [19] Quantitative Studies: How Many Users to Test? <http://www.nngroup.com/articles/quantitative-studies-how-many-users/>.
- [20] The OAuth 2.0 Authorization Framework. <http://tools.ietf.org/html/rfc6749>.
- [21] What is OpenID? <http://openid.net/get-an-openid/what-is-openid/>.