An Investigation into a Usable Identity Binding Service

Tristan Lewis
The MITRE Corporation
202 Burlington Rd
Bedford, MA 01730
tmlewis@mitre.org

William Kim
The MITRE Corporation
202 Burlington Rd
Bedford, MA 01730
wkim@mitre.org

Jill L. Drury
The MITRE Corporation
202 Burlington Rd
Bedford, MA 01730
jldrury@mitre.org

ABSTRACT

Managing online identities has become an increasingly challenging task. Often users utilize multiple personal and professional accounts on the same service. Sustaining multiple accounts for the same user at an application can lead to problems in security, privacy, and usability. In this paper we propose a solution to allow users to easily support multiple accounts while still maintaining secure operations and preserving user privacy. Further, we present a plan for determining whether this Identity Binding Service will be usable in two dimensions: ease of use and ease of comprehending identity management relationships.

1. INTRODUCTION

1.1 Problem

Managing digital identities can be burdensome to users and applications (Alotaibi and Wald, 2013). With the advent of distributed authentication technologies on the web, there are now multiple identity providers across different security domains, causing identity overload and password fatigue (Jøsang et al., 2007). For instance, each user may have: a personal account linked to friends, a personal professional identity, and a work account linked to their current employment—and potentially multiple accounts within any of these domains.

Users may want to bind some or all of these accounts together to facilitate a connected view across multiple accounts. In addition, relying parties (that is, applications that use authentication and identity claims from identity providers) may want more complete knowledge of users' multiple identities. The challenge is how to do this in a way that preserves authentication context, privacy, ease of use, and comprehensibility.

Currently, relying parties (RPs) bind identities either via custom solutions, or via "identity bridges" that are effectively a "man in the middle" representing the user to the application. We believe that custom implementations are inherently non-standardized, and identity bridges offer a solution that is fragile and not usable. Both present challenges for scaling and interoperability.

The usability of identity management is important in at least two dimensions. Of course, it should be easy to use. Of at least equal importance is the need to make the identity binding relationships clear to users. Due to the privacy implications of user binding data, it is critical that users know what is happening when the binding occurs. Our research aims to address the identity management usability flaw known as "cognitive scalability" (Dhamija and Dusseault, 2008): the need to reduce total workload and mental overhead even when identity management tasks increase in scope and complexity.

1.2 Solution

We have been creating a prototype Identity Binding Service that will allow users to explicitly correlate multiple identities with each other and will be independent of the RP, identity provider, and user. This service will provide an "opt-in" method for binding these identities that does not lose the original authentication context of the log in. (Using a "man in the middle" bridge component loses this authentication context.) Our service will allow RPs to query what accounts are correlated and will provide users the opportunity to link some or all of their accounts.

In effect, our proposed solution makes the Identity Binding Service a third party that requires mutual trust by both the RP and the user. Because this solution is designed as an independent component, it will not interfere with the original login mechanisms between the RP and the identity providers. For example: when using OpenID Connect (OpenID, 2015) for authentication, users may specify what (and where) information is shared.

Our research aims to show that, when our Identity Binding Service is implemented in a usable way, users will have greater satisfaction and control over online representation of their digital identities. We hypothesize that if RPs can consume the binding information, they will know more about the users and provide them a better experience. At this stage of our research, we are concerned with the user identity binding aspect of this problem.

2. ARCHITECTURAL APPROACH

The Identity Binding Service operates in two major phases: user binding and querying binding information. The focus of this research is to ensure users understand what is occurring in both phases, especially during phase 1, when they bind their identities.

2.1 Binding By the User

The user logs into the Binding Service using one of the identity providers (step 1 in Figure 1). On a successful authentication, the identity provider passes the authentication context to the Identity Binding Service. The Identity Binding Service then asks if the user would like to bind any other accounts. If so, the service passes the user off to the requested identity provider to repeat the process (step 2). This approach relies on a mechanism to allow simultaneous logins from multiple identity providers at the Identity Binding Service. If there are any previously bound identities to any of the logged-in identities, binding will create a superset of all the bound identities.

2.2 Querying Binding Information

Once an RP successfully authenticates a user (steps 3 or 4), the application queries the Identity Binding Service to see if that user is known by any other identities (step 5). After a successful

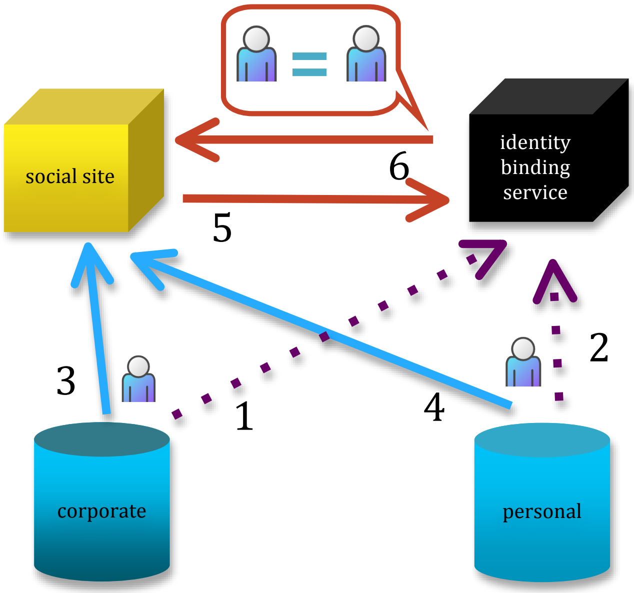

Figure 1: Relationship diagram between the User, Relying Party, and Identity Binding Service

authentication, the application queries the Identity Binding Service. In the simple case, the Identity Binding Service finds no other identities and the user is granted access to the application. If there are known bound identities, then the Identity Binding Service responds with those identities and the application can modify its response accordingly (step 6).

3. USABILITY TESTING PLANS

In addition to developing a prototype Identity Binding Service, we also developed a prototype RP that can utilize the binding service. These prototypes form a testbed to conduct user testing employing a between-subjects design, in which one half of the test participants will use the Identity Binding Service and the other half will log into the application using multiple accounts but without using the Identity Binding Service. To address the first usability dimension (ease of use), our testing measures will include task completion, number of mistakes, and overall satisfaction with the mechanism of logging into two separate accounts. They may also include measures from the Identity Access Management System (IAMS) Framework (Alotaibi and Wald, 2012).

Assessing the second usability dimension, which consists of comprehension of the identity binding process and relationships, requires more creative measures. We plan to elicit the user's comprehension of how the Identity Binding Service is working by asking participants to draw directed graphs showing the relationships among the RP, the Identity Binding Service, the identity providers, and themselves. Participants will be requested to label each edge with a number indicating its order in the execution sequence. Immediately after drawing the graphs, we will ask participants to verbally describe how they work, allowing us to audio-record participants' correct and mistaken perceptions.

This approach of eliciting mental models via drawing graphs is similar to that used by Sun et al. (2011, 2013), who asked experiment participants to draw how they believed log-in information flowed among entities in a specific scenario. Based on Sun et al.'s experience, we feel that our user group will be comfortable with drawing a simple graph. To ensure success, however, we will first guide participants through drawing a practice graph in a different domain such as using an automated teller machine.

Graph analysis will proceed by examining the similarity of the graphs to the correct representation as implemented in the prototype design. We will score the graphs based on the number of incorrect arcs, for example by using the similarity metrics of Zager and Verghese (2008). Further, the scores will be modified based on the number of arcs that are omitted, repeated unnecessarily, or executed out of order. The graph analysis will be augmented by a categorization of the users' mental models as extracted from the audio recordings.

We anticipate that changes to the interface will be needed based on the user study results. Once changes are made, we will execute another round of user tests.

4. ACKNOWLEDGMENTS

We would like to thank Justin Richer for his earlier and seminal contributions to this project. This work was supported by the MITRE Corporation Innovation Project 19MSR084-DA under the Information Technology Investment Area. © 2015 The MITRE Corporation. All Rights Reserved. Approved for Public Release; distribution unlimited; case #15-1599.

5. REFERENCES

Alotaibi, S. J., & Wald, M. (2012). IAMS Framework: A new framework for acceptable user experiences for integrating physical and virtual identity access management systems. In *Proc. of the World Congress on Internet Security (WorldCIS-2012).*

Alotaibi, S. J., & Wald, M. (2013). Evaluation of the UTAUT Model for Acceptable User Experiences in Identity Access Management Systems. In *Proc. of the Conference for the Internet Technology and Secured Transactions*. IEEE.

Dhamija, R., & Dusseault, L. (2008). The Seven Flaws of Identity Management: Usability and Security Challenges. *Security & Privacy , 6* (2), 24-29.

Jøsang, A., Zomai, M. A., & Suriadi, S. (2007). Usability and privacy in identity management architectures. In *Proceedings of the fifth Australasian Symposium on ACSW frontiers.*

OpenID Foundation (2015, April). OpenID Connect Core 1.0.

Sun, S.-T., Pospisil, E., Muslukhov, I., Dindar, N., Hawkey, K., & Beznosov, K. (2013, November). Investigating Users' Perspectives of Web Single Sign-On: Conceptual Gaps and Acceptance Model. *ACM Transactions on Internet Technology, 13* (1).

Sun, S.-T., Pospisil, E., Muslukhov, I., Dindar, N., Hawkey, K., & Beznosov, K. (2011). What Makes Users Refuse Web Single Sign-On? An Empirical Investigation of OpenID. *Symposium on Usable Privacy and Security (SOUPS)*. Pittsburgh: ACM.

Zager, L. A. & Verghese, G. C. (2008). Graph similarity scoring and matching. *Applied Mathematics Letters*, 21(1), 86-94.