Open Problems in Web 2.0 User Content Sharing

San-Tsai Sun and Konstantin Beznosov

University of British Columbia
Vancouver, Canada
{santsais,beznosov}@ece.ubc.ca

Abstract. Users need useful mechanisms for sharing their Web 2.0 content with each other in a controlled manner across boundaries of content-hosting and service providers (CSPs). In this paper, we discuss open problems and research opportunities in the domain of Web 2.0 content sharing among users. We explore issues in the categories of user needs, current sharing solutions provided by CSPs, and distributed access-control related technologies. For each open problem, we discuss existing and potential solutions, and point out areas for future work.

1 Introduction

"Web 2.0" is not necessarily the next version of Web technologies. Rather, it is, among other things, the concept of users being both consumers and producers of web content [Ore07]. Examples of user content are blogs, wikis, documents (such as Google Docs), multimedia (i.e. pictures, videos, music), personal biographical information, calendars, addresses, user's past and present physical location information, people the user is "linked to", history of visited pages and search queries, and social network groups that the user is a member of. One salient Web 2.0 characteristic important for this discussion is the ability of users without special technical skills to generate and share content on the Web.

More and more sensitive content such as political and social opinions, predilections and bias, biographical facts, and other person-identifiable content is being produced and shared by the users of the Web. What's more, the sensitivity of content may vary with time and context. For example, a recent high-school graduate might find it entertaining to share with the whole world pictures of him being silly at his graduation party. Four years later, these same pictures could become somewhat embarrassing, and even possibly a career obstacle for the same person when he is looking for a job after college. Yet, 15 years later, the pictures might become a subject of acute nostalgia for his youth. As such, he might want to share them with close friends and family. It is commonly believed [Bis05] that the sharing and dissemination of user content needs to be controlled in a way useful for its owners in order to unlock the full potential of user content production. Here are just few examples of Web 2.0 content sharing:

- Fiona, a FaceBook user, wants to share her photos on FaceBook with Bob, who does not have an account with FaceBook. Bob, in turn, decides to share this data with Charlie. How can Fiona do such sharing and yet retain control over who Bob shares her photos with, for how long, and how much of it?
- Sharing a shopping or todo list among family members.
- A job applicant creates a portfolio comprising his linkedin.com profile, professional essays in the form of blog posts, and his publications lists from ACM, IEEE, and DBLP. He wants to share the portfolio with a prospective employer but not with anybody else.
- Bob donates his medical record information in semi-anonymized form through Google Health for research purposes. But he wants to make sure that it's used only by medical researchers and only for the purpose of research.

For the purpose of illustrating our discussion, we will use the following scenario of content sharing as a running example:

Scenario 1 *Alice is a Girl Scout in the Colonial Coast Adventures (CCA) club. She and other CCA Girl Scouts took pictures at a scout training event, and would like to use their favorite photo web sites to share those photos online. Alice chooses MyPhoto.com. In CCA, pictures of training events can only be seen by CCA troop members and their parents, and Alice would like to implement this policy. She wants to limit access accordingly and is willing to share with all troop members and their parents. Jenny is another CCA member and Mary is her mother. In order for Mary to access Alice's photos for this event, Mary has to prove that she is the parent of Jenny and that Jenny is a CCA member. However, neither Jenny nor Mary are registered members of MyPhoto.com, and Alice does not know Mary.*

To summarize the overarching technical problem, *Web 2.0 users without special technical skills need useful mechanisms for sharing their content with each other in a controlled manner across content-hosting or application service provider (CSP) boundaries.* The corresponding research problem in Web 2.0 content sharing is the lack of understanding of *the factors that influence the design* of useful mechanisms for controlled content sharing among users across CSPs.

The rest of the paper elaborates on specific technical challenges—and the corresponding research opportunities—that need to be addressed in order to enable secure Web content sharing. In Section 2, we discuss the sharing and control needs of Web 2.0 users. Once the user needs are understood better, a general Web 2.0 user-centric content sharing model is presented in Section 3 to frame the technical aspects of our discussion. Section 4 investigates the current solutions provided by CSPs, and Section 5 discusses open problems related to distributed access-control technologies. For each open problem, we discuss potential solutions and point out areas for future research. At the end, Section 6 summarizes and concludes the paper.

2 User Needs

Problem 1 [**Sharing Needs**] *What are the users' needs for sharing their content?*

Discussion:

Before investigating a solution for controlled sharing, the very first step has to be toward understanding the users' needs for sharing. What do the users need content sharing for? Given that content sharing is a secondary task, what are the corresponding primary tasks and objectives? What are the scenarios of sharing content and data?

What kind of produced content and personal data do they want to share? In what circumstance do the users share content? Only when they are sitting in front of their descktops/laptops at home and office, or on the go, using their mobile devices? For how long do they want sharing to persist: few days, weeks, or months?

With whom do they want to share? Just friends, family, and colleagues? Are there any direct correspondences between the sharing groups on the Web and social groups in the real world? Can the sharing groups be described using some digital or physical attributes, or they can only be enumerated? How often does the membership in these groups change? How does the change in membership affect sharing decisions?

How do the users prefer sharing to be done, and why? How often do they perform the task of sharing? What factors affect the preferences for sharing and its frequency? What produced content and personal data users don't want to share at all and why?

To explore preferences for general information sharing, Olson et al. [OGH05] investigated what content is shared and with whom. They found that participants abstract the details of sharing into high-level classes of recipients and information which are treated similarly. Voida et al. [VEN+06] studied the sharing practices of ten employees at a medium-size research organization to identify the types of content they share and with whom, the mechanisms they use to share, and how much control over the information they grant to the sharing recipients. They identified 34 different types of files that are shared among colleagues, friends, and families. In their results, email is the most common mechanism for sharing (45%), followed by network folders (16%) and posting content to a web site (11%). The study also identified the breakdowns that users have experienced in their file sharing activities. The main classes of breakdowns are (1) difficulties in selecting a sharing mechanism with desired features that are available to all sharing participants, (2) forgetting what files had been shared and with whom, and (3) problems in knowing when new content was made available. Similarly, Whalen [Wha08] investigated file sharing practice in both work and personal context of about 200 employees at a US research institution. Most of her results confirm the findings made by Voida et al. In addition, she identifies the factors that influence the choice of sharing method used. She also found that lacking ac-

tivity information (e.g., who and when) on the shared files could be problematic for both security and collaboration.

The above user studies provide preliminary knowledge about sharing practices with respect to files residing on a user's computer or workspace. For Web content sharing, Miller et al. [ME07] conducted an empirical study of photo sharing practices on Flickr.com. They found that privacy-concerned users primarily use e-mail, supplemented with Web galleries, to control the privacy level of their photos. The perception of using email for sharing by those users is that an e-mail message is intentional, requires no setup, and is targeted at a specific list of recipients. Miller et al. suggested that a sharing solution should look and feel much like e-mail, but with a more robust underlying framework geared towards photo sharing.

Although Web content sharing is a common practice, relatively few studies exist in research literature investigating this topic. Further user studies on understanding Web content sharing practices and the issues users encounter are important research topics that should be investigated.

Problem 2 [**Control Needs**] *What are the users' needs for controlling their content sharing?*

Discussion:

Studying the needs of the users for sharing their content and data is a necessary but not a sufficient step towards solving the general problem of Web 2.0 content sharing. What kind of control Web users need is another important practical and research question.

First of all, what content do they want to share without control, and why? Since publicly accessible content constitutes the majority of Web data, how do the users decide which content they want to control?

For the content and data they want to restrict access to, what control granularity do they need? Here are just some examples: album vs. picture/vedio in album, blog vs. thread vs. post, calendar vs. events in calendar.

What time factors are important for sharing control? Possible options are: (1) period, e.g., for next 3 weeks, (2) date, e.g., until/after May 5. It could very well be that the users do not need to restrict sharing based on time. Before time-sensitive controls are pursued further, the users' need for them needs to be determined.

Grouping is a widely used technique for scaling administration and run-time overhead. It has been used since the days of Bell-LaPadula (BLP) [BL73] information flow policies, and later in Role-Based Access Control [SCFY96] (RBAC), and by many other approaches to scale authorization. At the same time, grouping abstractions can be highly counterintuitive for users. CORBA authorization architecture[YD96], for instance, scales nicely due to the use of policy domains, as well as required and granted rights for grouping objects and their methods. However, the architecture makes it hard to express even simple role-based policies [BD99]. How do the users want to group shared content and those with whom they share it? Some options are: attribute-based groups (e.g., all colleagues with

@ece.ubc.ca e-mail address, all undergraduate students from the academic program, all those with whom Alice already shared these pictures), delegation of group membership management (e.g., my family and their friends), grouping of the shared content. How should the management of those group memberships be done? Can group memberships be inferred automatically? For example, can Alice define group *girlscouts* of CCA as her *friends*?

Given the high diversity of Web 2.0 users, how usable should the control be? Input and output capabilities of mobile devices will likely affect usability. What other factors affect the usability of controls, and how?

The interplay of factors such as granularity, scalability, and usability is likely to become entangled with one another. This makes it necessary to study all three together or find a way to dissect the "knot." Recent studies of employing both restricted natural language [BKK06] for expressing organizational privacy policies and GUI for administering Windows ACLs [MR05,RBC+08] suggest that this is a tricky business, which can lead to unexpected outcomes. On the other hand, making the state of the controls opaque to the end users can lead to dangerously inaccurate mental models of the controls [RHB09].

The need for right revocation and the ways of supporting it are other subjects to be investigated. Given that Web content can always be copied by authorized users, the benefit of revocation might be moot. What are, if any, the users' needs for revocation?

Do users need to have uniform controls across CSPs? If so, how can it be made usable, given the heterogeneity of the content. If not, how shall the controls differ?

Problem 3 [**Demographic Factors**] *How do answers to the above questions vary with the demographics of users?*

Discussion:

How do independent variables (e.g., gender, age, education, occupation) affect sharing needs, preferences, patterns, and other aspects identified above?

Investigation of user needs and demographic factors requires field work using qualitative methods, such as interviews [Kva96], contextual enquiry [BH98], naturalistic observations, and other ethnographic instruments [Fet98]. More quantitative methods, such as surveys, should help validating qualitative findings and generalize them to larger populations of users.

3 General Model of Web 2.0 Content Sharing

Once the user needs are understood better, investigating the capabilities of existing and new technical solutions becomes a viable next step. To ground our discussion in concrete terms, we frame it in the language of a general Web 2.0 content sharing model as illustrated in Figure 1. In this model, a *user* is not only a content *owner* and *consumer*, but a *credential* issuer as well. A user enrolls a set of *identities* (e.g., user name/password) with multiple *identity providers* to represent themselves when accessing shared content and constructing access

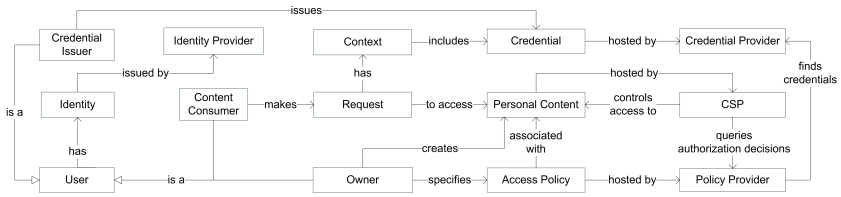

Fig. 1. Web 2.0 user-centric content sharing model.

polices. A content owner creates personal content on a CSP and associates that content with access-control polices that are hosted by her *policy provider*. To access shared content, a content consumer chooses an appropriate identity to initiate a *request*. Each request contains the identity provided by the consumer and a corresponding set of *context* information. Context information is the meta-data of a request, such as user-specific profile attributes, current location, date/time of the request and user *credentials*. A credential is an assertion of a certified user-attribute from another individual user or an organization authority. To mediate access, a CSP requests authorization decisions from a policy provider to protect shared content. The policy provider then acts as a policy decision point (PDP) which responds with authorization decisions based on the context of request and a set of pre-defined credentials and access policies.

User-centric content sharing requires user-centric authentication and authorization mechanisms. For user-centric authentication, the user should be able to control her own identities and should be free to choose when and where to use them. For user-centric authorization, access decisions are based on the policies associated with the protected content and the related credentials issued by credential issuers. The content owner should have the freedom to choose policy providers to host policies, and trusted authorities to issue credentials. In a user-centric Web, an access policy follows the user. One access policy hosted by one policy provider should be able to be enforced across CSPs.

4 Current Sharing Solutions by CSPs

A secure Web 2.0 content sharing solution should be inter-operable across administrative domains and should be usable for average web users. In this section, we explore open problems related to sharing solutions currently provided by CSPs.

Problem 4 [**Existing Sharing Solutions**] *How well do the existing approaches support users' needs of controlled sharing?*

Discussion:

With Web 2.0, the user is both a consumer and provider of Web content [Ore07]. However, the current Web is site centric as opposed to the user centric Web model discussed in the previous section. Figure 2 illustrates the current site-centric Web. For each service provider, a Web user has to maintain a separate

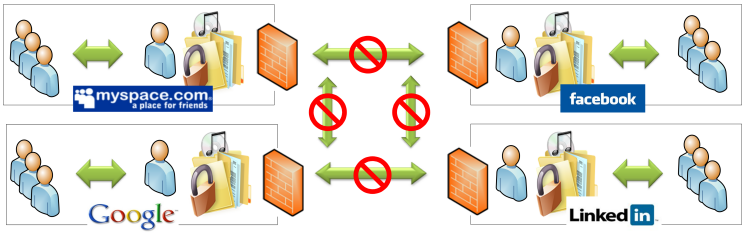

Fig. 2. Site-centric Web.

copy of identity, personal data, and relationships with other Web users as . Each CSP forms an administrative domain by defining the membership of users, permissions on protected resources, and access-control policies. In this paper, we use the term "*walled garden*" to refer to such an administrative domain defined by a service provider. Since each walled garden controls its own set of users and employs a different access-control mechanism to protect personal content, it is difficult to share personal content beyond walled gardens.

Personal content sharing is currently available in limited forms. There are three main content sharing mechanisms offered by content-hosting or application service providers. The first one is to make user content public. Obviously, this is inadequate for controlled sharing. The second one is a mechanism we labeled as the *walled garden* approach. With this approach, the user who "owns" content can grant permissions directly to other users (or user groups) within the same CSP. The walled-garden approach is easy to implement and use. Its main limitation is that not all the desired content users (e.g., Girl Scouts and their parents) are necessarily registered with the corresponding CSP; and thus, users outside of that CSP cannot be granted selective access. Even within the same walled garden, the resource requester and owner might not be known to each other (e.g., Alice does not know some other Girl Scouts and their parents who use MyPhoto.com), increasing the challenge of controlled sharing for both the owners and consumers of content.

To enable controlled sharing beyond walled gardens, some CSPs (e.g., Google, Facebook, Flickr) use a third mechanism, which we refer to as the *secret-link* approach. A secret-link is a hard-to-guess URL that uniquely identifies a shared resource (e.g., http://spreadsheets.google.com/ccc?key=px0Ox4z1SIE). When a resource owner shares personal content using a secret-link, the corresponding CSP creates a special URL for that resource. Anyone who knows a secret-link can access the content referred by that link. To share specific personal content, a resource owner sends (sometimes with the aid of the CSP) the secret link via email to select users. The message recipients view the shared content by clicking on the link. Secret-link is easy to use for both owners and users, and it provides a certain degree of control over sharing since only those who obtained (or guessed) the link can access the content. However, the use of secret-link has limitations. The main one is that the secret-link can be forwarded or otherwise

"leaked" to unauthorized users. In other words, if sharing content using secret-links is viewed as a capability-based access control, then unrestricted capability delegation becomes an issue. Some CSPs, such as Google Picasa, provide a "sign-in required to view" option when sending a secret-link. However, this requires the recipients to have an account with the CSP in order to view the shared content, essentially reducing the approach to a walled garden approach.

Since the walled-garden and secret link approaches have the advantage of wide deployment, their support for the user needs deserves further investigation.

5 Distributed Access-Control Technologies

Federated identity systems allow reuse of a user identity across multiple CSPs, and distributed authorization systems provide expressive policy languages for expressing access policies with sound algorithms for deriving authorization decisions. Those technologies could potentially be used or extended to provide solutions to the problem of Web 2.0 content sharing. In this section, we discuss open problems related to federal identity management and distributed authorization systems.

Problem 5 [**Federated Identity**] *How can the approaches to federated identity be improved to better support controlled sharing across walled gardens?*

Discussion:

One of the limitations of the walled-garden approach for sharing is the challenge of having separate user identifiers with different CSPs. Users end up maintaining multiple identities and corresponding passwords at multiple sites, which leads to weaker passwords and/or password re-use across accounts [FH07].

Federated identity management enables user attributes in one CSP to be provided to other CSPs in a federation. Solutions such as coalition-based access control (CBAC) [CTWS02], Liberty Alliance Project [Lib02], Shibboleth [Int08] (based on SAML [OAS02]), and OpenID [RF07] are examples of federated identity systems. Federated identity solutions enable cross-domain single sign-on, and remove the need for users to keep identifiers and passwords at individual CSPs. Identity providers (IdPs) supply user information, while CSPs "consume" provided identity and mediate accesses based on this information. Except OpenID, current federated identity solutions require pre-established relationships between IdPs and CSPs in a federation, essentially forming a larger walled garden by aggregating existing ones.

OpenID is an open protocol for identity management and web-based single-sign-on. OpenID uses a URI as an end-user's identifier, which acts as a universal user account and is valid across all CSPs. OpenID is user-centric in the sense that users are free to choose their own OpenID identity providers. However, OpenID suffers from adoption and usability problems, is vulnerable to phishing attacks, and creates privacy problems.

Currently, there are many OpenID IdPs, but not enough CSPs that support OpenID. The "identity war" started since the beginning of the Web, and "walls"

have been built by CSPs to protect their subscriber base. Even before OpenID, major CSPs have been providing a way (e.g., Microsoft Windows Live ID, Yahoo BBAuth [Yah08], AOL OpenAuth [AOL08], Google AuthSub [Goo08]) for other CSPs to accept user credentials from their domain. However, most CSPs would like to maintain as many registered users as possible and are reluctant to tear down their guarded walls. Even though Facebook has recently created a precedent of accepting OpenID accounts [Ion0p], it remains to be seen if the "identity war" will end anytime soon.

Phishing attacks on the federated identity protocols are a looming threat. OpenID and other similar protocols (e.g., Google AuthSub [Goo08], AOL OpenAuth [AOL08], Yahoo BBAuth [Yah08]) may cause users to become accustomed to being redirected to identity provider websites for authentication. If users do not verify the authenticity of these websites before entering their credentials (and they usually do not [SDOF07,DTH06]), phishing attacks are possible. An attacker can send an email to lure users to a phony service provider and redirect users to a forged site where they are asked to enter their passwords. To prevent phishing attacks, users must confirm the authenticity of an identity provider before entering their credentials. Existing research on authenticating web-sites to users include security indicator [Fra05,Cor05,Net08,HJ08], secure bookmark with known websites [DT05,WML06,YS06,PKP06], and automated detection and blacklisting of known phishing sites [Ear08,Net08]. However, studies suggest that security indicators are ineffective at preventing phishing attacks [DTH06,SDOF07], and blacklisting known phishing sites still suffers from high rate of false-positives and false-negatives [ZECH07]. Even with improved security indicators, users still tend to ignore them [WMG06,SDOF07].

Tracking of users through their "global" identifiers raises a privacy concern. OpenID protocol—as defined by its specification [RF07]—enables IdPs to track all websites a user logged into using her OpenID account. The tracking capability of OpenID makes cross-site profiling easy and possible. How to prevent IdPs from tracking websites their users have visited is an open question.

The usability of federated identity schemes could be improved. OpenID uses a URI as a user's identifier, but Web users perceive a URI as a "web address" instead of a personal identifier. In the context of Web content sharing, users rarely know the URIs of those with whom they want to share their content with, but they tend to know each other's e-mail addresses. Using an email address as user account is common, and a user is often prompted for an email address as user name. For instance, major service providers (e.g., Google, Yahoo, AOL) use email addresses as user accounts to associate provided services. Thus, user experience for registering and entering CSPs that support OpenID_{email} would be the same as they experience today. Extending federated identity protocols (such as OpenID) to use email addresses as alternative identifiers might improve their usefulness for Web users.

Problem 6 [**Distributed Authorization Systems**] *Whether and how distributed authorization systems can be employed to enable controlled content sharing among users?*

Discussion:

Traditional access control mechanisms make authorizations based on the identity of the request (e.g., access-control matrix, RBAC[SCFY96]). However, in decentralized environments such as the Web, the content owner and the requestor often are unknown to each other (e.g., Alice does not know Mary and Jenny). There is a substantial body of literature addressing the problem of authorization within distributed environments.

PolicyMaker [BFL96] coined the term "trust management" to denote an access control model in which authorization decisions are based on locally stored security policies and distributed credentials (signed statements), without explicit authentication of requestor's identity and a centralized repository of access rules. Policies and credentials in PolicyMaker consist of programs written in a general programming language such as AWK. Although general, it is very hard to understand the overall access policy for a protected resource. KeyNote [BFIK99], the next version of PolicyMaker, uses a C-like notation and regular expression syntax for describing conditions. SPKI/SDSI [EFL$^+$99] is a digital certificate scheme for authorization, which provides methods for binding authorization privileges to keys and for localized name spaces and linked local names. A credential in KeyNote and SPKI/SDSI delegates certain permissions from an issuer to a subject. A chain of credentials can be viewed as a capability which authorizes the subject at the end of the chain. RT [LMW02] is a family of languages that add the notion of RBAC to the concept of trust management systems such as KeyNote and SPKI/SDSI. Akenti [TJM$^+$99] is a user-centric access-control system using X.509 attribute certificates. XACML [XT05] (eXtensible Access Control Markup Language) defines a generic authorization architecture and a policy language for expressing and exchanging access policy using XML.

ABLP Logic [LABW91,ABLP93], proposed by Abadi, Burrows, Lampson, Plotkin, is a propositional modal logic for authentication in distributed systems. ABLP Logic can be used to determine who is the original requestor that made a request which has been through multiple channels and processes on distributed hosts. Higher-order logic allows application-specific modal logic to be defined. Appel et al. [AF99] demonstrated how ABLP Logic can be expressed as application-specific definitions and lemmas using AF Logic—a general higher-order logic with only a few inference rules. Based on the framework proposed by Appel et al., the user requesting access is responsible for constructing a proof, and the server simply checks that proof— a concept called proof-carrying authorization (PCA). Using the concept of PCA, Bauer et al. implemented an access-control system for regulating access to web pages [BSF02].

Relationships between user and owner are intuitive to Web users and are commonly used to derive authorization decisions. Most walled garden sharing mechanisms provided by CSPs support a certain degree of relationship-based access control. Carminati [CFP06] et al. proposed an access control mechanism for web-based social networks, where policies are expressed as constraints on the type, depth, and trust level of existing social relationships. Lockr [TGS$^+$08] is another access control mechanism based on social relationships. In Lockr, a

policy (social access control list) is a static set of required social relationships, and a credential (social attestations) is a signed statement stating the social relationship between issuer and receiver.

Existing distributed authorization systems provide expressive policy languages for expressing credentials and access policies, sound algorithms for discovering credential chains, and inference engines for deriving authorization decisions. However, because Web 2.0 access policies for personal content are authored by users without special technical skills, and are enforced by mutual-untrusted walled gardens, there are many remaining issues to address. The main issues are *usability* and *inter-operability*. The expressive power of a policy language must be balanced with usability. An average internet user must be able to comprehend the language to ensure that an access policy matches the owner's sharing intention. For sharing content on the Web, the system should not require any special software to be installed. To share personal content beyond walled gardens, credential and access policy that are authored in one policy provider should be employable to protect personal content residing on multiple CSPs. The data owner should have the freedom to choose policy providers, and the access polices should follow the owner to wherever she goes. In addition to usability and inter-operability, *granularity* of control and *accountability* should be considered as well. Content created by Web users is diverse and sometimes complex; the content owner should be able to specify access-control in a fine-grained format. For example, owners might want to protect a photo in an album, an event in a calendar, or even a paragraph within a blog. For accountability, the owner should be able to know which data is being accessed, by who and when, and be able to revoke an authorization at anytime if necessary.

Thus, the research questions need to be investigate include: Are those mechanisms usable for average Web users? How can access policies/credentials expressed in one system be employed and enforced by multiple CSPs? How common vocabularies in credential and policy statements can be agreed upon amongst Web users? How distributed credential-chain discovery processes can be carried out among CSPs? What degree of control granularity do those mechanisms provide? What activity information on shared content should be visible to content owners?

6 Summary

In this paper, we explore open problems and research opportunities in the domain of Web 2.0 content sharing. For each open problem, we discuss potential solutions, and point out areas for future research. We do not claim the problems we have identified represent a complete list. While many issues may exist, we only pick problems that are directly related to design and implementation phases of a Web 2.0 content sharing solution. Other research questions that are related to adoption and evaluation phases (e.g., what should be a success criteria for a solution in the space of Web 2.0 controlled sharing? How to compare solutions? How to evaluate a solution?) might need to be investigated as well. Since secure

content sharing in Web 2.0 is a complex problem, how to partition the problem into smaller and more manageable parts is also an important question.

We believe that user-centric access control is a fundamental part of a user-centric Web. In the user-centric Web, the user is in charge. Users own their personal content and are free to share it across walled gardens. In the user-centric Web, users also have the freedom to choose their favorite providers for their identities, content, social relationships, and access-control policies. The separation of personal content and services puts the focus of a service provider on providing valuable services to the user it serves, forcing the service provider to be just a service provider—no longer an identity or social graph hogger.

Acknowledgements

This paper is a result of several discussions with LERSSE members. Special thanks go to Kirstie Hawkey and Matei Ripeanu of LERSSE for their detailed input on the earlier presentations of the points discussed in this paper.

References

ABLP93. M. Abadi, M. Burrows, B. Lampson, and G. Plotkin. A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems*, 15(4):706–734, 1993.

AF99. Andrew W. Appel and Edward W. Felten. Proof-carrying authentication. In *Proceedings of the 6th ACM conference on Computer and communications security*, pages 52–62, New York, NY, USA, 1999.

AOL08. AOL LLC. AOL Open Authentication API (OpenAuth). http://dev.aol.com/api/openauth, January 2008.

BD99. Konstantin Beznosov and Yi Deng. A framework for implementing role-based access control using CORBA security service. In *Fourth ACM Workshop on Role-Based Access Control*, pages 19–30, Fairfax, Virginia, USA, 1999.

BFIK99. Matt Blaze, Joan Feigenbaum, John Ioannidis, and Angelos D. Keromytis. The KeyNote trust-management system version 2, September 1999.

BFL96. Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *the 1996 IEEE Symposium on Security and Privacy*, pages 164–173, Washington DC, USA, 1996.

BH98. Hugh Beyer and Karen Holtzblatt. *Contextual Design, Defining Customer-Centered Systems*. Morgan Kaufmann Publishers, San Francisco, CA, 1998.

Bis05. Matt Bishop. Psychological acceptability revisited. In Lorrie Faith Cranor and Simson Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, chapter 1, pages 1–11. O'Reilly Media, Inc., 2005.

BKK06. Carolyn A. Brodie, Clare-Marie Karat, and John Karat. An empirical study of natural language parsing of privacy policy rules using the sparcle policy workbench. In *SOUPS '06: Proceedings of the second symposium on Usable privacy and security*, pages 8–19, New York, NY, USA, 2006. ACM.

BL73. D.E. Bell and L. LaPadula. Secure computer systems: Mathematical foundations. Technical Report MTR-2547, Volume I, Mitre Corporation, Bedford, Massachusetts, 1973.

BSF02. Lujo Bauer, Michael A. Schneider, and Edward W. Felten. A general and flexible access-control system for the web. In *Proceedings of the 11th USENIX Security Symposium*, pages 93–108, Berkeley, CA, USA, 2002. USENIX Association.

CFP06. Barbara Carminati, Elena Ferrari, and Andrea Perego. Rule-based access control for social networks. In *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops*. LNCS, Springer-Verlag, 2006.

Cor05. CoreStreet Ltd. Spoofstick, 2005.

CTWS02. Eve Cohen, Roshan K. Thomas, William Winsborough, and Deborah Shands. Models for coalition-based access control (cbac). In *Proceedings of the seventh ACM symposium on access control models and technologies*, pages 97–106, Monterey, California, USA, 2002.

DT05. Rachna Dhamija and J. D. Tygar. The battle against phishing: Dynamic security skins. In *SOUPS '05: Proceedings of the 2005 symposium on Usable privacy and security*, pages 77–88, New York, NY, USA, 2005. ACM.

DTH06. Rachna Dhamija, J. D. Tygar, and Marti Hearst. Why phishing works. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 581–590, New York, NY, USA, 2006. ACM.

Ear08. Earthlink Inc. Earthlink toolbar: scambloker for windows users, 2008.

EFL+99. Carl M. Ellison, Bill Frantz, Butler Lampson, Ron Rivest, Brian Thomas, and Tatu Ylonen. SPKI certificate theory, September 1999.

Fet98. David M Fetterman. *Ethnography: Step by Step*. Sage Publications Inc., 1998.

FH07. Dinei Florencio and Cormac Herley. A large-scale study of web password habits. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 657–666, New York, NY, USA, 2007. ACM.

Fra05. Rob Franco. Better website identification and extended validation certificates in ie7 and other browsers, 2005.

Goo08. Google Inc. Authsub authentication for web applications. http://code.google.com/apis/accounts/docs/AuthSub.html, December 2008.

HJ08. Amir Herzberg and Ahmad Jbara. Security and identification indicators for browsers against spoofing and phishing attacks. *ACM Trans. Interet Technology.*, 8(4):1–36, 2008.

Int08. Internet2. Shibboleth System. http://shibboleth.internet2.edu/, 2008.

Ion0p. Daniel Ionescu. Facebook Embraces OpenID. http://www.pcworld.com/article/165110/, 200p.

Kva96. S. Kvale. *InterViews: An Introduction to Qualitative Research Interviewing*. Sage Publications, 1996.

LABW91. Butler Lampson, Martin Abadi, Michael Burrows, and Edward Wobber. Authentication in distributed systems: Theory and practics. In *ACM Symposium on Operating Systems Principles*, pages 165–182, Asilomar Conference Center, Pacific Grove, California, 1991.

Lib02. Liberty Alliance. Liberty Alliance Project. http://www.projectliberty.org/, 2002.

LMW02. Ninghui Li, John C. Mitchell, and William H. Winsborough. Design of a role-based trust-management framework. In *SP '02: Proceedings of the 2002 IEEE Symposium on Security and Privacy*, page 114, 2002.

ME07. Andrew D. Miller and W. Keith Edwards. Give and take: A study of consumer photo-sharing culture and practice. In *Proceedings of the CHI 2007*, pages 347–356, San Jose, California, USA, April 28 –May 3 2007.

MR05. Roy A. Maxion and Robert W. Reeder. Improving user-interface dependability through mitigation of human error. *International Journal of Human-Computer Studies*, 63:25–50, 2005.

Net08. Netcraft Ltd. Netcraft toolbar, 2008.

OAS02. OASIS. Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML). http://www.oasis-open.org/committees/security/docs/cs-sstc-core-00.pdf, April 2002.

OGH05. Judith S. Olson, Jonathan Grudin, and Eric Horvitz. A study of preferences for sharing and privacy. In *CHI '05 extended abstracts on Human factors in computing systems (CHI '05)*, pages 1985–1988, New York, NY, USA, 2005. ACM.

Ore07. Tim Oreilly. What is Web 2.0: Design patterns and business models for the next generation of software. *Communications and Strategies, No. 1, p. 17*, 2007.

PKP06. Bryan Parno, Cynthia Kuo, and Adrian Perrig. Phoolproof phishing prevention. In *Proceedings of the 10th International Conference on Financial Cryptography and Data Security*, volume 4107, pages 1–19. LNCS, February 27 2006.

RBC+08. Robert W. Reeder, Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, Kelli Bacon, Keisha How, and Heather Strong. Expandable grids for visualizing and authoring computer security policies. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1473–1482, New York, NY, USA, 2008. ACM.

RF07. David Recordon and Brad Fitzpatrick. OpenID authentication 2.0 - final. http://openid.net/specs/openid-authentication-2_0.html, December 2007.

RHB09. Fahimeh Raja, Kirstie Hawkey, and Konstantin Beznosov. Towards improving mental models of personal firewall users. In *CHI '09 extended abstracts on Human factors in computing systems*, page 6, Boston, MA, USA, April 2009. ACM.

SCFY96. Ravi Sandhu, Edward Coyne, Hal Feinstein, and Charles Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, 1996.

SDOF07. Stuart E. Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. The emperor's new security indicators. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, pages 51–65, Washington, DC, USA, 2007. IEEE Computer Society.

TGS+08. Amin Tootoonchian, Kiran K. Gollu, Stefan Saroiu, Yashar Ganjali, and Alec Wolman. Lockr: social access control for web 2.0. In *Proceedings of the first workshop on Online social networks*, pages 43–48, Seattle, WA, USA, 2008.

TJM+99. Mary Thompson, William Johnston, Srilekha Mudumbai, Gary Hoo, Keith Jackson, and Abdelilah Essiari. Certificate-based access control for widely distributed resources. In *Proceedings of the 8th USENIX Security Symposium*, pages 215–228, Washington, D.C., USA, August 23–26 1999.

VEN+06. Stephen Voida, W. Keith Edwards, Mark W. Newman, Rebecca E. Grinter, and Nicolas Ducheneaut. Share and share alike: exploring the user interface affordances of file sharing. In *Proceedings of the SIGCHI conference on Human Factors in computing systems CHI '06:*, pages 221–230, New York, NY, USA, 2006. ACM.

Wha08. Tara Whalen. *Supporting file sharing through improved awareness.* Ph.D. dissertation, Dalhousie University, Canada, 2008.

WMG06. Min Wu, Robert C. Miller, and Simson L. Garfinkel. Do security toolbars actually prevent phishing attacks? In *Proceedings of the SIGCHI conference on Human Factors in computing systems(CHI '06)*, pages 601–610, New York, NY, USA, 2006. ACM.

WML06. Min Wu, Robert C. Miller, and Greg Little. Web wallet: preventing phishing attacks by revealing user intentions. In *SOUPS '06: Proceedings of the second symposium on Usable privacy and security*, pages 102–113, New York, NY, USA, 2006. ACM.

XT05. XACML-TC. OASIS eXtensible Access Control Markup Language (XACML) version 2.0. OASIS Standard, 1 February 2005.

Yah08. Yahoo Inc. Browser-Based Authentication (BBAuth). http://developer.yahoo.com/auth/, December 2008.

YD96. Zhonghua Yang and Keith Duddy. CORBA: a platform for distributed object computing. *SIGOPS Oper. Syst. Rev.*, 30(2):4–31, 1996.

YS06. Ka-Ping Yee and Kragen Sitaker. Passpet: convenient password management and phishing protection. In *SOUPS '06: Proceedings of the second symposium on Usable privacy and security*, pages 32–43, New York, NY, USA, 2006. ACM.

ZECH07. Yue Zhang, Serge Egelma, Lorrie Cranor, and Jason Hong. Phinding phish: Evaluating anti-phishing tools. In *Proceedings of the 14th Annual Network and Distibuted System Security Symposium (NDSS 2007)*, 2007.