Policy Derived Access Rights in the Social Cloud

Ferry Hendrikx and Kris Bubendorfer
School of Engineering and
Computer Science
Victoria University of Wellington
New Zealand
Email: ferry,kris@ecs.vuw.ac.nz

***Abstract*—Social clouds are a relatively new paradigm that allow users of an underlying social network to share their resources with their "friends", using previously established relationships. However, this sharing has a number of issues, including granularity of friendships, resource costs and maintenance. In this paper we argue that sharing decisions should be based on relationship information augmented by supplementary metadata derived from multiple sources. Users should be able to leverage the information available on their non-uniform friend relationships when making decisions, allowing them to confidently share their resources with those that would normally be outside of their immediate social circle. We introduce GRAft, our Generalised Recommendation Architecture, that provides us with a mechanism to support this new approach.**

I. Introduction

Since the launch of the first recognizable social network SixDegrees in 1997 [1], social networks have seen rapid evolution and massive growth. For example, Facebook has grown to over 1 billion users, of which 655 million are daily users[1]. These sites allow family, friends and colleagues to stay in touch, discover events and share digital resources. The network of "friends" that are created on these social networks are utilised by Social Clouds [2] to share resources such as information, services and hardware. Social clouds are open and created in an ad-hoc fashion, with resources being contributed and dynamically shared [3]. There are a number of issues that may be identified in these systems:

- **Identity**. There is no guarantee of the real-world identity of any user [4], [5]. There have been some "hacks" involving the cloning of entire personal profiles in order to convince others of a fake identity[2].
- **Granularity**. There may not necessarily be a close friend relationship between users wanting to share resources. Until recently, the granularity of relationships on most social networks was too coarse, allowing only the simplest boolean representation.
- **Cost**. A user may limit the sharing of their resources because of the cost, or inherent risk of providing these resources to others.
- **Risk**. A user may not fully trust a resource because of the risks associated with its use (such as availability, accuracy, dependability, integrity [6], etc).
- **Maintenance**. A user needs to maintain a list of their "friends" and the rights or permissions associated with each. This becomes progressively more time-consuming as the user gathers more "friends" and as the complexity of the relationships increase.

The granularity issue was partially addressed by Facebook in 2012, when it added the notion of "close friends" [4], while the frendship relationship in Google+[3] is somewhat more complete in that it allows users to place their friends into different "circles", emphasising the different types of relationship. Granularity can also be resolved using classic access control approaches. Access control is a process that requires "every access to a system and its resources be controlled and that all and only authorised accesses can take place" [7]. Access control policies can be divided into three groups: Discretionary, Mandatory and Role-Based [7], [8]. However, the application of classic access control approaches curtails the ease and usefulness of sharing in a social cloud. For example, the sharing of a resource with a research group might involve searching for all potential users in the group and then creating a "friend" relationship to each one, before finally assigning access permissions. Further, regular maintenance would be required to add, update and remove users.

In this paper, we propose a mechanism that addresses the issues identified above by compensating for the lack of information available on a "friend". Multiple existing sources of information are leveraged to provide us with supplementary information about others. Sharing with friends and people outside of our immediate social circle, with whom we do not have an existing direct relationship, can be enabled by examining other's attributes to identify those with whom we can share. In our approach, a user's behaviour and demographics are evaluated by policies and converted into access permissions and roles. The information about a user that is used to derive access control is obtained from multiple places, including databases, directories and non-obvious sources such as forums, wikis and other social media.

The lack of a "close friend" relationship (or finer) and whether all friends should be treated the same way are both addressed by appropriate policies. Similarly, the cost and risk issues can also be addressed using policies, although risk could be further addressed by treating resources like services and evaluating their past "behaviour". Finally, user maintenance issues are addressed by the decoupling of policy and the data collection mechanisms.

[1] https://newsroom.fb.com/Key-Facts - last accessed June 2013
[2] http://www.snopes.com/computer/facebook/pirates.asp
[3] https://plus.google.com/

978-0-7695-5083-1/13 $25.00 © 2013 IEEE
DOI 10.1109/eScience.2013.27

The rest of the paper is organised as follows. In the next section we provide an overview of the social cloud followed by section III in which we introduce social cloud policies. We then present related work in Section V, and conclude the paper in Section VI.

II. Social Cloud Overview

A. Cloud Computing

Cloud computing has gained widespread attention in the last few years, in part due to its ability to leverage economies of scale to deliver cost-effective infrastructure and software as services. Its growth has been driven largely by increasing expectations of computing power, storage and ubiquitous access to the internet [9].

The United States' National Institute of Standards and Technology has provided a useful definition for cloud computing: "*Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction*"[4].

The cloud computing paradigm allows for resource utilization via low-level abstractions [2]. The services provided by low-level cloud service providers such as Amazon (EC2 and S3), Google (App Engine) and Microsoft (Azure) are often found in the offerings of high-level cloud service providers. For example, the cloud-based file synchronization and backup service offered by Dropbox[5] uses Amazon's S3 service for its storage.

B. Social Networks

In parallel to cloud computing, the growth of social networks and their associated use means that large numbers of people now use this medium every day, with almost 50% of users checking it 3 times a day or more [10]. Boyd defines social networks as "*web-based services that allow individuals to (1) construct a public or semi-public profile within a bounded system, (2) articulate a list of other users with whom they share a connection, and (3) view and traverse their list of connections and those made by others within the system*" [1]. Of these three points, the second is key, as it allows a user to nominate those other users that he or she knows. Interestingly, these connections between users often represent real-world ties, allowing users to communicate with others that are already part of their extended real-world social network [1].

C. Social Cloud

These established trust relationships between members of a social network can be utilized by a social cloud, allowing users to easily share their resources (such as storage and services) with their "friends". In [4], [11], a social cloud is defined as a "resource and service sharing framework utilizing relationships and policies established between members of a social network". In effect, a social cloud utilizes the underlying social network relationships to control and manage access to cloud computing resources.

In [4], the authors introduce a number of application scenarios that derive their key benefits from being in a social cloud. Common features of these application scenarios include decentralized infrastructure and management, and the utilization of social relationships to drive sharing of resources:

- Social Computation cloud. The ability for users to easily share their under-utilized computing hardware with others (such as friends and charities).
- Social Storage cloud. The ability for users to share, backup and replicate their data. In particular, photos are often already shared with friends, removing some of the security maintenance issues commonly encountered with data sharing.
- Social Collaborative cloud. Allows collaborations to share resources using social networks. This lowers barriers in the creation of new communities and facilitates simpler sharing.
- Social Science cloud. The ability for users to contribute their resources towards collaborative science problems that have captured community interest. Some command examples include SETI@Home, Folding@Home and BOINC.

In each of these scenarios, both the user sharing and the user consuming resources need the ability to select who they will interact with, and to what degree. This is accomplished using a set of preferences that a user may have established before any sharing takes place. Selecting a list of users that you may want to share with is trivial, however defining the access rights for each and every user in that list is potentially a much larger task. This problem is exacerbated when you consider that some studies have shown mean numbers of "friends" on Facebook exceeding 240 people [12].

One approach to this problem, shown in Figure 1, is to utilise a socio-technical adapter [13] that enhances the "friend" relationship data that is held in the social network with supplementary metadata. This additional information allows a user to compensate for their non-uniform relationships, and still maintain effective control via policies. For example, in the case described above, a policy can be used to control sharing, obviating the need to define the access rights for every "friend". Our Generalised Recommendation Architecture (GRAft) in the case of the Social Cloud acts as a distributed socio-technical adaptor that collects recommendation information from multiple sources and makes it available for use in consumer applications. A potential key usage of this information is in the evaluation of access control policies.

III. Social Cloud Access Policies

In this section we present three sample social policies, implemented within GRAft using Ruler[6]. For a scenerio such as a F2F photo storage cloud, an example GRAft policy that expresses that storage is accessible to *Family and the most interactive and closest of Friends* is:

[4]http://www.nist.gov/itl/cloud/

[5]http://www.dropbox.com/

[6]https://github.com/bobthecow/Ruler

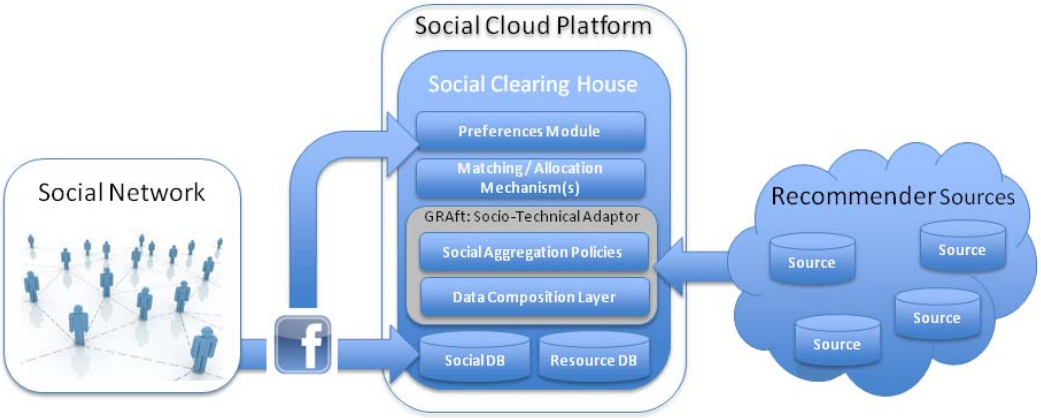

Fig. 1. The GRAft socio-technical adapter.

```
$rb->logicalOr(
    $rb['in_circle']->equalTo('family'),
    $rb->logicalAnd(
        $rb['degree_of_friendship']->lessThan(2),
        $rb['interactions']->greaterThan(100)
    )
)
```

In a more professional collaborative environment, the first example policy represents *Co-authors in the 'Distributed Systems' research group*, while the second policy represents *Co-authors, and Co-authors of co-authors, that are employed by either Acme Corporation or Studentville University*.

```
$rb->logicalAnd(
    $rb['degree_of_coauthorship']->lessThan(2),
    $rb['research_group']->equalTo('Distributed␣
        Systems')
)

$rb->logicalAnd(
    $rb['degree_of_coauthorship']->lessThan(3),
    $rb->logicalOr(
        $rb['employer']->equalTo('Acme␣Corporation')
            ,
        $rb['employer']->equalTo('Studentville␣
            University')
    )
)
```

IV. GRAft Overview

GRAft is a distributed and open infrastructure that collects and stores recommendation information about participants. The collected information may then be used and examined by other participants. The recommendation information stored in GRAft may include such things as reputation, competency and demographics. The participants may be either users or services, and are all identified using OpenID [14].

There are multiple sources of information in GRAft. Each source pushes updated information about its users into the GRAft network. Consumer applications obtain the information they require from GRAft, and use it as "input" when evaluating their access control policies. More detail on GRAft is available in [15].

V. Related Work

A social cloud depends largely on the trust relationships established between its users. In [16], the authors discuss the usage of friendship and co-authorship social graphs as input to their social cloud. In [17], the authors introduce the foundations for the contextualization of trust within a social cloud. They state that using existing relationships between users is more efficient than a relationship established between anonymous individuals. This idea is expanded in [13], where the concept of a socio-technical adapter is introduced. The GRAft work presented in this paper provides an implementation of a socio-technical adapter.

In [3], the authors introduce a framework for the sharing of resources in a social cloud for the scientific community. Their work is built on Facebook, and allows for the sharing of computational resources through Virtual Machines (VM). Access to each VM is controlled by membership to social groups within Facebook. The work addresses access restrictions in only a limited way.

A system that allows the sharing of information using social networks is presented in [18]. A social network is calculated for every user. Users are able to attach access control lists (ACLs) to their resources, and for every access, the system decides if the current user may have access. Although this work has similarities to ours, the access model is explicit, in that an ACL must be built by the owner for every resource they want to make available. Access to all resources is boolean and inflexible.

In [19], the authors introduce myExperiment, a workflow discovery and sharing system. Workflows are created and shared publicly or via contacts in a social network. Users may download and enact workflows, and upload modified versions that can similarley be shared with others. Although the sharing of workflows via social networks has similarities to our work, myExperiment does not have the ability to construct arbitrary policies using social metrics.

Pythia, a reputation-based authorization system is introduced in [20], [21]. A central repository stores reputation information obtained from multiple applications. Relying parties

may then query the system to obtain reputation information about users that has been processed through a rules engine. Pythia differs from our work in three distinct ways: it has a centralised architecture and is not distributed. Pythia provides relying parties with reputation information that has been processed by the system, whereas GRAft makes available all of the key information, and leaves any calculation to the relying party. Lastly, Pythia is focussed on reputation-based models, and does not consider other recommendation types.

A Cross-Community Reputation (CCR) model is introduced in [22]. The key arguments made about the CCR model are that reputation information obtained from multiple sources is more accurate, and that this approach obviates the need to bootstrap a new reputation for a user when it interacts with a new community. The model includes reputation sharing and conversion, and allows for the mapping of information from one source to another. A separate policy controls how much information may be shared across the different communities. This work differs from our own in that it only discusses the sharing of reputation information, and does not provide an architecture or implementation details.

VI. Conclusion

Social clouds allow "friends" to share their resources in an open and dynamic fashion. However, this sharing has a number of issues, including granularity of friendships, resouce cost and maintenance that make it less than ideal to share resources openly. Classic access controls could be implemented to facilitate sharing of a user's resources, however this does not address all of the issues, and removes some of the usefulness and ease of sharing in a social cloud.

In our approach, metadata about others is captured and converted by policies into access permissions and roles. In essence, a user's relationship to the owner of a resource, and supplementary information about that user are considered when granting access to a resource. The transformation from this metadata into access controls addresses some of the issues identified in social cloud sharing, and introduces flexibility in the form of policies. The policies are implemented using metrics and allow for both novel and malleable access control over resources that would otherwise not be easily available.

References

[1] D. M. Boyd and N. B. Ellison, "Social Network Sites: Definition, History, and Scholarship," *Journal of Computer-Mediated Communication*, vol. 13, no. 1, pp. 210–230, 2007.

[2] K. Chard, S. Caton, O. Rana, and K. Bubendorfer, "Social Cloud: Cloud Computing in Social Networks," in *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*. IEEE, 2010, pp. 99–106.

[3] A. M. Thaufeeg, K. Bubendorfer, and K. Chard, "Collaborative eResearch in a Social Cloud," in *E-Science (e-Science), 2011 IEEE 7th International Conference on*. IEEE, 2011, pp. 224–231.

[4] K. Chard, K. Bubendorfer, S. Caton, and O. Rana, "Social Cloud Computing: A Vision for Socially Motivated Resource Sharing," *Services Computing, IEEE Transactions on*, vol. 5, no. 4, pp. 551–563, 2012.

[5] B. Chen and A. Roscoe, "Social networks for importing and exporting security," in *Large-Scale Complex IT Systems. Development, Operation and Management*. Springer, 2012, pp. 132–147.

[6] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka, and J. Molina, "Controlling data in the cloud: outsourcing computation without outsourcing control," in *Proceedings of the 2009 ACM workshop on Cloud computing security*. ACM, 2009, pp. 85–90.

[7] P. Samarati and S. C. de Vimercati, "Access control: Policies, models, and mechanisms," in *Foundations of Security Analysis and Design*. Springer, 2001, pp. 137–196.

[8] R. S. Sandhu and P. Samarati, "Access control: principle and practice," *Communications Magazine, IEEE*, vol. 32, no. 9, pp. 40–48, 1994.

[9] G. Pallis, "Cloud computing: The new frontier of internet computing," *Internet Computing, IEEE*, vol. 14, no. 5, pp. 70–73, 2010.

[10] B. Debatin, J. P. Lovejoy, A.-K. Horn, and B. N. Hughes, "Facebook and online privacy: Attitudes, behaviors, and unintended consequences," *Journal of Computer-Mediated Communication*, vol. 15, no. 1, pp. 83–108, 2009.

[11] K. John, K. Bubendorfer, and K. Chard, "A Social Cloud for Public eResearch," in *E-Science (e-Science), 2011 IEEE 7th International Conference on*. IEEE, 2011, pp. 363–370.

[12] J. B. Walther, B. Van Der Heide, S.-Y. Kim, D. Westerman, and S. T. Tong, "The Role of Friends' Appearance and Behavior on Evaluations of Individuals on Facebook: Are We Known by the Company We Keep?" *Human Communication Research*, vol. 34, no. 1, pp. 28–49, 2008.

[13] S. Caton, C. Haas, K. Chard, K. Bubendorfer, and O. Rana, "A Social Compute Cloud: Allocating and Sharing Infrastructure Resources via Social Networks," *Submitted to IEEE Transactions on Services Computing*, vol. ?, no. ?, pp. ?–?, 2013.

[14] D. Recordon and D. Reed, "OpenID 2.0: a platform for user-centric identity management," in *Proceedings of the second ACM workshop on Digital identity management*, ser. DIM '06. New York, NY, USA: ACM, 2006, pp. 11–16. [Online]. Available: http://doi.acm.org/10.1145/1179529.1179532

[15] F. Hendrikx and K. Bubendorfer, "Malleable access rights to establish and enable scientific collaboration," in *Proceedings of the 9th IEEE International Conference on eScience*. Beijing, China: IEEE, October 2013.

[16] A. Mohaisen, H. Tran, A. Chandra, and Y. Kim, "Socialcloud: Using social networks for building distributed computing services," *arXiv preprint arXiv:1112.2254*, 2011.

[17] S. Caton, C. Dukat, T. Grenz, C. Haas, M. Pfadenhauer, and C. Weinhardt, "Foundations of trust: Contextualising trust in social clouds," in *Cloud and Green Computing (CGC), 2012 Second International Conference on*. IEEE, 2012, pp. 424–429.

[18] J. Mori, T. Sugiyama, and Y. Matsuo, "Real-world oriented information sharing using social networks," in *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*. ACM, 2005, pp. 81–84.

[19] D. De Roure, C. Goble, and R. Stevens, "The design and realisation of the myExperiment virtual research environment for social sharing of workflows," *Future Generation Computer Systems*, vol. 25, pp. 561–567, 2009.

[20] P. J. Windley, D. Daley, B. Cutler, and K. Tew, "Using reputation to augment explicit authorization," in *Proceedings of the 2007 ACM workshop on Digital identity management*, ser. DIM '07. New York, NY, USA: ACM, 2007, pp. 72–81. [Online]. Available: http://doi.acm.org/10.1145/1314403.1314416

[21] P. J. Windley, K. Tew, and D. Daley, "A framework for building reputation systems," *WWW 2007*, pp. 8–12, 2007.

[22] T. Grinshpoun, N. Gal-Oz, A. Meisels, and E. Gudes, "CCR: A model for sharing reputation knowledge across virtual communities," in *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 01*, ser. WI-IAT '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 34–41. [Online]. Available: http://dx.doi.org/10.1109/WI-IAT.2009.13