Privacy-by-design in Federated Social Web Applications

Alexandre Passant, Owen Sacco and Julia Anaya
Digital Enterprise Research Institute (DERI)
National University of Ireland, Galway
Galway,Ireland
firstname.lastname@deri.org

ABSTRACT

In this paper, we propose a privacy-by-design architecture for Federated Social Web applications, in particular, focusing on federated microblogging environments. This architecture is applied to SMOB. a Semantic Microblogging framework, which will be extended to use a distributed broadcasting architecture such as Google's PubSubHubbub and also it will incorporate privacy preferences for filtering the distribution of posts.

1. PRIVACY IN FEDERATED SOCIAL WEB APPLICATIONS

As discussed in the recent report by the W3C Social Web Incubator Group [3], there is a need for a "Standards-based, Open and Privacy-aware Social Web". Various initiatives are currently following this path, such as Diaspora, GNU Social, identi.ca or SMOB. Yet, while most of them focus on the "open" and "standard-based" features, the privacy aspect must still be tackled properly.

Our vision of the Federated Social Web shall allow everyone to decide whom they want to share information with thanks to simple settings that any user can understand – even if back-ended by complex policies or rule languages. In addition, we want to go further and restrict for instance republication of confidential information, but ensuring that privacy policies are associated with every piece of content published on the Web, so that they remain associated with it when the content moves from one place to another.

In this paper, we discuss our recent findings on the topic, and our contribution that consists in a privacy-by-design architecture for Federated Social Web applications, with a special focus on the microblogging context done in the context of a Google Research Award. In particular, we discuss how we combine (1) a lightweight privacy preference ontology [2] together with (2) a distributed broadcasting architecture for micro-content and (3) tag-based policies, the whole stack being build open Google's PubSubHubbub and using state-of-the-art Semantic Web standards (SPARQL and SPARQL Update, etc.). Using this model, one could decide to restrict all content tagged with #websci to colleagues only, while content tagged #mylife will be only broadcasted to local friends – and unavailable to people that are not identified as colleagues or friends.

Of particular relevance is the way our policies (and group membership) are interpreted in the cloud, so that the users allowed to access the content are identified on runtime, based on dynamic policies. For instance, one would define "colleague = listed in DERI's FOAF profile", and that policy will be interpreted in the cloud when the message is sent. That way, the content broadcaster is able to identify who can access the content at a particular time, without the need for the user to specify a static list of people to send the information to, also simplifying the task from an end-user perspective.

We will demonstrate the approach in the context of SMOB [1], our distributed microblogging platform, and will provide some formative evaluation (based on user- surveys) and concrete figures regarding the scalability of the approach, that could be seamlessly integrated in other social web context since it is based on open Web (and Semantic Web) standards.

2. REFERENCES

[1] A. Passant, J. Breslin, and S. Decker. Rethinking Microblogging: Open, Distributed, Semantic. In *Proceedings of the 10th International Conference on Web Engineering, ICWE'10*, 2010.

[2] O. Sacco and A. Passant. A Privacy Preference Ontology (PPO) for Linked Data. In *Proceedings of the Linked Data on the Web Workshop, LDOW2011*, 2011.

[3] W3C Social Web Incubator Group. A Standards-based, Open and Privacy-aware Social Web. Technical report, 2010.