Report
W3C – OpenSocial Foundation workshop "Social Standards: The Future of Business"
7-8 August 2013, San Francisco, California

Executive Summary

The W3C-OpenSocial Foundation joint workshop on "Social Standards: The Future of Business" convened with the goal to formulate a standardization strategy to make "social" a first-class citizen of the Web. The current landscape around social applications on the Web is fragmented, holding back business results. The Open Web Platform, with W3C's royalty-free patent policy, offer a sound base for moving forward. The workshop --, hosted by AppFusions, and sponsored by IBM and the Open Mobile Alliance --, featured two-days of presentations and intense discussions of challenges, use cases, and potential standardization strategies for social. The workshop brought together over 70 people from over 50 companies around this theme, ranging from start-ups like Crushpath to established enterprise social networks like Yammer.

A number of points of consensus emerged:

- The OpenSocial Foundation and W3C will co-ordinate to establish one or more new Working Group(s) to create the next version of OpenSocial on top of the Open Web Platform

- The common data-format for social activities should be JSON-based ActivityStreams, and a new version that supports extensible data formats be pursued.

- As the "back-end" of social sites usually involves loosely coupled "NoSQL" graph-based data, work on property graphs should be pursued in a Community Group.

The workshop built a strong alliance between the OpenSocial Foundation and W3C to make "social" a first-class citizen of the Web. Given the move to mobile and device-independence (or multi-device dependence), building a common framework for social Web applications on top of HTML5 is a top priority. Leveraging the complementary strengths of both organizations can help make social standards a success.

This workshop built on the work of the Social Business Community Group, in particular the Social Business Jam in 2011 and reports. Previous workshops on social included the Federated Social Web conference and the Future of Social Networking Workshop. The W3C feels that thanks to the work of the Social Business Community Group, critical mass has been reached within W3C membership to make social standards a full-featured part of the Open Web Platform.

Meeting minutes for August 7th and August 8th are available online.

The Business Case for Social Standards

Dion Hinchcliffe (Dachis Group) opened the workshop with a keynote calling social media the largest communication revolution since the Web itself. The lack of underlying standards to social media stunts businesses' ability to get in touch with their own users and customers. Simplicity is key to success with standards for business, Hinchcliffe said, pointing to the success of RSS over Web Services. With "hundreds of social networks with over a million users" unable to talk to each other, the network effect is being lost.

Mark Crawford (SAP) then explained, based on SAP's experience with SuccessFactors, that social needs to move from personal relationships based on proprietary stacks to "business" relationships based on standards so that social activities (learning, internal communities, HR management, supply chains, onboarding new employees) can be integrated into business processes rather than lost in the "black hole" of email. Since companies employ many diverse packages of software across heterogeneous environments, standards are the way forward.

Ed Krebs (Ford) followed by detailing a reference architecture that showed both how fragmented the current landscape was and also pointed to the possibilities for a unified architecture to make it easy for engineers to build enterprise social networks that can successfully interoperate both within and between enterprises such as Ford. Storing files in multiple places is not efficient, and "nuggets of wisdom" are lost that are crucial to the business.

Don Buddenbaum (IBM) presented on how 'social' has to be embedded where people do their work, with metrics included so businesses can understand the concrete results of using social.

Lloyd Fassett (Azteria) gave a presentation on how social standards could enable businesses to move from "pipes to platforms" that enable a business to make better choices in use-cases such as health-care staffing.

Use-cases

What concrete use-cases could be addressed using social standards?

Li Ding (Memect) provided an analysis suggesting that all use-cases could be thought in terms of providing an extended memory for a business.

Monica Wilkinson (Crushpath) discussed how standards help start-ups ship working code faster, saying that her start-up deploys a vast variety of standards, ranging from de-facto closed work such as the Open Graph Protocol (Facebook's "Like" Button) to community-driven work like ActivityStreams.

Eric Meeks (University of San Francisco) argued that Linked Data complements OpenSocial, and demonstrated how it enables academic social networking.

Adam Boyet (Boeing) pointed out how their custom-built InSite social platform allows "connections everywhere" to enable both internal and external collaboration and expert-finding for Boeing, but authentication and federation of identity profiles were major pain-points in integrating InSite with other products like Microsoft SharePoint.

Lastly, **Dan Schutzer** (FSTC) reminded the audience that deployment in the financial sector, depended on privacy and security to protect users, as well as a focus on risk and compliance that are necessary to deal with anti-fraud and disclosure requirements.

A focus on expert-finding, as well as identifying the right context for expertise, was heavily discussed.

Discussion took place over the difference between the emergent proposed social platform and traditional collaboration software, with the key difference being that collaboration software was focused on pre-existing teams while social software was meant to help a business discover connections that it might not even have known existed before -- both between the employees of a business and between a business and its customers.

Social Standards Architecture

Monica Lam's keynote on "How Mobile Revolutionizes Social" raised the case that mobile could revolutionize social, as phones are essentially thin clients to social networks. Lam followed with a demonstration of an application that let users create their own ad-hoc social networks without servers based on their phone. As social standards are currently spread across multiple standards-bodies and

grass-roots efforts, how can we unite them into a coherent "social platform" built on top of the Open Web Platform?

Bryan Sullivan (AT&T) noted that a social architecture would have to scale globally in a mobile environment, and demonstrated how the Open Mobile Alliance had already constructed a draft architecture (SNEW) based on pre-existing work such as OStatus. There was still much work to be done, such as integration with NFC and multi-factor authentication, and ActivityStreams templates were needed to standardize various common workflows. Lastly, user control and privacy were still major open issues.

A mobile social networking could even increase network efficiencies, and **Fabio Mondin** (Telecom Italia) demonstrated how their work with the eCousing project allowed reduced network usage by, for example, placing social content closer to the location of the event. To enable these kinds of use-cases, the social networking architecture needs to be able to communicate with the networking architecture.

Jason Gary (IBM) pointed out how events and roles need to be embedded in ActivityStreams, but currently profiles do not support roles despite roles having the ability to be the "killer app" for social. Discussion over the importance of roles between roles was brought up.

Ashok Malhotra (Oracle) brought up the fact that the back-end of social networking sites store massive property graphs, a graph-based data-structure where lists of properties are attached to each node. Currently the details are different for how each vendor stores property graphs, and Oracle would be willing to make a submission to the W3C to start work in standardizing them. It is currently unclear how much of property graphs could be handled by the RDF data model.

Federating the Social Web

Matt Franklin (W20 Digital) started with a call to action on how a new generation of standards to federate the social web, building on top of OpenSocial and ActivityStreams, would be necessary. In particular, OpenSocial does not address identity and the social graph, and ActivityStreams needs to have better interoperability with processing rules and levels of visibility.

Given that the proposed next version of ActivityStreams is using the JSON-LD format, **Gregg Kellogg** presented on how JSON-LD adds URIs and links to JSON, thus making JSON compatible with the RDF data model.

Ed Krebs (Ford) presented that any federated architecture needs to have a "PubSubHub" system are needed such that new business systems can feed data to each other without changing the other servers.

Sam Goto (Google) presented on how schema.org was being extended to take on actions (essentially a taxonomy of verbs), similar to the "Embedded Experiences" of OpenSocial where verbs can take on well-defined subjects and objects with semantic roles.

Theodoros Michalareas (VELTI) presented on the OPENi API, which after reviewing over 140 APIs to produced, using principles of privacy-by-design, an API for federated identities and app-produced contexts.

There was considerable discussion over the choice of data-formats (HTML with Microformats2, JSON-LD, ordinary JSON) as well as the relationship of context to security concerns.

Next Steps for OpenSocial

OpenSocial is the foremost API for enterprise social applications, and its evolution will help drive the open social web. In the OpenSocial "State of the Union" address, **Mark Weitzel** (Jive) and **Andy Smith** (IBM) laid out a plan for building the next version of OpenSocial on top of the Open Web Platform. OpenSocial has always been focused on securely sharing context bi-directionally with applications. A new version of OpenSocial that builds on top of Shadow DOM and Web Components will let developers build OpenSocial applications in the same style as any other HTML5 application while maintaining OpenSocial's ability to share context and create "embedded experience" that prevent users from losing their context.

Building on their points, **Beth Lavender** (MITRE) discussed how their work allowed MITRE to view a business either at a particular point in time or view the business as activities were occurring in "real-time."

The host of the workshop, **Ellen Feaheny** (AppFusions), discussed how AppFusions makes standards like OAuth talk to each other in their rapid integration of Jive, IBM, and Atlassian applications.

Shane Caraveo (Mozilla) presented Mozilla's new "Social API" that embeds capabilities to the user agent's sidebar such as notifications, social bookmarking, share, and chat windows.

Dimitri Glazkov (Google) then gave an in-depth presentation on Web Components, which led to considerable excitement on how OpenSocial could work together with Web Components and other new capabilities being developed in HTML5.

Running Code

Inspired by the "IndieWeb Camps" and "Federated Social Web Summits", the workshop hosted a session of demonstrations of running code. **Tantek Çelik** began by introducing the idea of "IndieWeb", based on the twin principles of Own your own data, Eat your own dogfood, and Publish Own Site, Syndicate Elsewhere.

Aaron Pareki showed how by running his own domain he could be his own identity server (IndieAuth), and then with **Bret Comnes** a demonstratedion was done showing how a watch could be used to authenticate into a site @@.

Evan Prodromou (Status.Net) presented his new "Pump.io" codebase for an ActivityStreams server with varying degrees of privacy, allowing streams to be filtered and writable only to certain groups.

Ben Werdmueller (Lakatoo) presented Idno, a social publishing platform built just on top of HTML5 and microformats. Users should be put first, and beware of putting technology before usability.

Patrick Deegan (ID3) demonstrated Open Mustard Seed that uses virtual machines to created trusted applications bundles. Their goal is to create a new social ecosystem of trusted digital institutions based on personal data. Access control (distribution control of ActivityStreams) and consumption of ActivityStreams were mentioned as outstanding problems, with a client API for ActivityStreams and WebMention brought up as possible solutions.

Next Steps

At the end of the workshop, break-out groups met to discuss areas to be standardized next. Groups formed around the following topics:

- OpenSocial and Gadgets will focus on radical simplification leveraging HTML5, moving from the XML definition of a gadget to a situation where AJAX requests are performed directly against a

page. How context works with cross-origin requests and how application tags can be supported by HTML5 are the next steps.

- ActivityStreams will focus on a new version, ActivityStreams 2.0, to increase extensibility and handle state. There was a large discussion over the role of JSON-LD as a syntax for ActivityStreams, but as ActivityStreams 2.0 does not depend on it, it was viewed as acceptable.

- Identity and Profile Federation needs to focus on a set of core attributes that show how previous work in the area (vCard, Microformats, PortableContacts) can be extended with desired features such as skill-levels and certifications. H It is necessary to understand how profiles federate using protocols such as Pubsubhubbub is necessary to understand.

- IndieWeb will focus on user experience, in particular making it much easier to use the reply button and work with browsers to make it easier to share content.

- Property Graphs need to have their data model defined, as well as APIs and schemas. Potential cross-over work on exploiting property graphs with the OpenSocial API should be investigated.

- Linked Data and vocabularies need to focus on how to create new kinds of vocabularies that can enable social business, such as expertise vocabularies.

Additionally **R.V. Guha** (Google) came to answer questions about licensing and transparency of schema.org. Guha noted that data a company marks up using schema.org microdata still belongs to the website, and so that data cannot be re-used without that website's permission, but that he would investigate whether changes to the schema.org terms of use were warranted.

Interest in following through with each of above topics was fairly well distributed, with more than ten people interested in continuing concrete work on each.

- The idea of a high-level "Social Business Architecture" document showing how all the diverse pieces could be put together in a use-case driven architecture also attracted significant interest.
- New W3C Working Groups on ActivityStreams, OpenSocial, and possibly federation should be pursued.
- Property Graphs and Profile work should happen in Community Groups in order to reach more maturity.
- Schema.org would continue to work with W3C and other grassroots communities to make its process more open and transparent for vocabularies.
- The Social Business Community Group would evolve to handle messaging and co-ordination responsibilities as an Interest Group.

Conclusion

All participants are invited to join the Social Business Community Group in order to help draft the charters for new work. Even if you missed the workshop, you can join the conversation to build the next version of social on top of the Web!

Comment [A1]: I underlined Working Groups in this bullet, and Community Groups in the next line, to emphasize the difference that is suggested.

Comment [A2]: Is "federation" here the same or different from what is suggested for "profile work"? If the same, then they should be combined somehow.

Comment [A3]: should? will? ("would" doesn't sound like correct English, but also could be stronger

Comment [A4]: should? may? (ditto to comment A3)