Final Report on Verifiable Claims by Industry Participants and Invited Experts

Web Payments Interest Group / Verifiable Claims Task Force

1. Background on the Verifiable Claims Task Force

In 2015, the Web Payments Interest Group had identified what seemed to be a need for verifiable claims for Web Payments Use Cases and launched the Verifiable Claims Task Force to:

1. Clarify if the **Problem Statement** *(below)* was accurate and an unsolved problem.
2. Gather data to determine if there was a desire to create an interoperable ecosystem around the expression and transmission of verifiable claims.
3. Seek expert advice among the people that were involved in the creation of LDAP, SAML, OpenID Connect, XDI, SSL/TLS, and attribute-exchange related technologies in general.

2. Problem Statement

There is currently no widely used ***user-centric*** and ***privacy-enhancing*** *(see Derived Requirements from Research)* standard for expressing and transacting verifiable claims (aka: credentials, attestations) via the Web.

These problems exist today:

- In existing attribute exchange architectures (like SAML, OpenID Connect, Login with SuperProviderX, etc.), users do not independently exist from service providers. This means users can't easily change their service provider without losing their digital identity. This leads to vendor lock-in, identity fragility, reduced competition in the marketplace, and reduced privacy for all stakeholders.
- There is no interoperable standard capable of expressing and transmitting rich verifiable claims that cuts across industries (e.g., finance, retail, education, and healthcare). This leads to industry-specific solutions that are costly, inefficient, proprietary, and inhibit users' ability to manage their digital identities in a cohesive way.
- There is no standard that makes it easy for users to assert their qualifications to a service provider (I am over the age of 21, I have an account at Bank X, my loyalty card number is Y, I am a citizen of the USA and here is my Passport, I am a Chartered Financial Analyst, etc.).

3. Summary of Research Findings

- There is broad consensus on the accuracy of the problem statement. There is no user-centric and privacy-enhancing ecosystem for issuing, storing, and requesting verifiable claims over the Web.
- There are clear use cases that have been documented in payments, insurance, banking, finance, healthcare, and education.
- There is broad consensus that current technologies are not readily addressing the problem.
- There is broad consensus that the W3C and IETF are capable of addressing the problem.
- There is broad consensus to re-use technologies and standards that exist today and modify them to address the problem statement, if possible. There is rough consensus that some new technologies may need to be created to do so.
- There is no consensus around which technologies should be used for each piece in the ecosystem; it is assumed that an officially chartered Working Group would determine the technologies and best mode.
- There is rough consensus that a minimum first step is to establish a way to express verifiable claims by either re-using existing technology or taking technology that has been incubated and standardizing it. There is rough consensus that perhaps both approaches may need to be executed in parallel.
- There is rough consensus in two groups on a verifiable claim issuing, storage, and request *protocol*. One group believes that starting on a verifiable claim issuing, storage, and request *protocol* is necessary, but may be difficult to get W3C member buy-in on. The other group believes that having a verifiable claim *data model and syntax* without a *protocol* would not address the problem statement.

4. Requirements Identified by Research Findings

For the purposes of this report, a **user-centric**, **privacy-enhancing**, verifiable claims ecosystem is defined as having the following requirements:

- The subject (e.g. a user) of a verifiable claim has independent existence; users can control and own their identifiers.
- Verifiable claims are a mechanism for transferring the trust issuers have in users to third parties; the trust model does not directly involve software agents or services.
- Users may ask issuers to make verifiable claims about them. Users can be given these claims in digital form, store them at a location of their choosing, and control with whom they are shared.
- Users may decide how to aggregate claims and manage their own digital identities.
- Users may use software agents to help them store claims, but these agents are not party to the claims; users may freely change agents without losing or affecting the meaning of their claims.
- Users may share verifiable claims without revealing the intended recipient to the software agent they use to store the claims.
- Consumers of verifiable claims may independently verify them without revealing their identity to the issuers of the claims.
- Users may be guaranteed basic unlinkability when sharing pseudo-anonymous verifiable claims with multiple parties.
- Semantics are decentralized. Vocabularies for verifiable claims may evolve independently in the communities that use them. Issuers may use these to indicate the meaning of their verifiable claims.
- Trust is decentralized. Consumers of verifiable claims decide which issuers to trust.

Table of Contents

5. Research Input Sources

5.1 Survey

Forty-three (43) organizations that either issue verifiable claims or consume verifiable claims were surveyed to determine their needs in this space. The organizations were:

The Bill and Melinda Gates Foundation, Accreditrust, Target Corporation, Dublin Core Metadata Initiative (DCMI), Verisys Corporation, New Zealand Government, Digital Bazaar, Walmart Stores, IMS Global Learning Consortium, The Hypothes.is Project, Education Testing Service, NACS, VU University Amsterdam, Knovation, Electronic Transactions Association, Pearson, Rabobank, VitalSource Technologies, Inc., CWI, Badge Alliance, Ingenico, D2L Corporation, The Dutch National Bank, Capella University, ConsenSys, The Paciello Group (TPG), Conexxus, Deque Systems, Citrix, The Open Group, Worldpay, Deutsche Telekom AG, INRIA, The Open Group, Raiseworth Pty Ltd, Ripple, GNU, MediaGoblin, Independent (ex-banker), Bloomberg, Focafet, BPCE, Federal Reserve Bank of Chicago, and OpenLink Software.

The aggregated results of the survey can be found here:

http://opencreds.org/presentations/2015/w3c-tpac/anonymized.html

5.2 Interviews

Expert interviews were performed with people that have been participating in the identity/verifiable claims space for over a decade. All interviews were performed by people in their personal capacity, not as a representative of their organization. Titles and organizations for each person are listed only as background for people that do not know who these people are (in no particular order):

- Brad Hill (Facebook, Co-Chair of W3C Web Application Security Working Group, ex-PayPal). Transcript and audio of interview.
- Christopher Allen (co-author of SSL 3.0 and TLS, Organizer of Rebooting Web of Trust Workshops). Transcript and audio of interview.
- Drummond Reed (Respect Network, Founding member of OpenID work, Founder of OpenID Exchange, Past Executive Director of Infocard Foundation). Transcript and audio of interview.
- John Tibbetts (Chief Product Architect, IMS Global Consortium (standards body for education sector), core architect Learning Tools Interoperability v2 Standard). Transcript of email interview.

- Jeanne Kitchens (Credential Transparency Initiative, Associate Director Center for Workforce Development at Southern Illinois University). Transcript and audio of interview.
- Bob Sheets (Credential Transparency Initiative, Research Professor at George Washington Institute for Public Policy). Transcript and audio of interview.
- David Chadwick (Professor of Information Systems Security at University of Kent). Transcript of email interview.
- Michael Schwartz (Chairman of OTTO Working Group, Kantara Initiative, implementer of GLU Framework including SAML, OpenID, OAuth, and other identity technologies, UMA Working Group). Transcript and audio of interview.
- Dick Hardt (Amazon, core participant in OpenID and OAuth work, well-known Identity 2.0 speaker). Transcript and audio of interview.
- Jeff Hodges (PayPal, co-author of LDAPv3, SAMLv1.x, SAMLv2, Liberty ID-FF (now SAMLv2), and ID-WSF). Transcript of email interview.
- Harry Halpin (W3C Staff Contact, IETF Liason). Transcript of email interview.
- David Singer (Apple). Transcript of email interview.

5.3 Credentials Community Group

The Credentials Community Group was created in August 2014 based on an identified need for verifiable claims in the Web Payments Community Group. The separate group was created to focus the discussion to payments in the Web Payments CG and verifiable claims/credentials in the Credentials CG. The community group consists of 77 participants that have been working on experimental technology and pre-standards specification documents that attempt to address the problem statement defined at the top of this report. It is expected that some of these Community Group specifications will be used as input to any future Working Group related to verifiable claims.

6. Areas of Consensus

The following sections outline areas of consensus.

6.1 No User-Centric, Privacy-Enhancing Ecosystem Exists for Verifiable Claims

Interviewees were asked if the Problem Statement defined on the first page of this report was accurate. Responses ranged from “It’s more or less okay” to “very strong agreement”. When there were disagreements with the problem statement, they tended to focus around a particular sentence (e.g. “It’s not quite accurate that no system exists for exchanging claims, SAML for instance kinda does this") or word (e.g. "user-centric" is problematic). As a result, the Problem Statement was adjusted slightly to address the concerns raised during the interviews. There was consensus around the problem statement that "No User-Centric, Privacy-Enhancing Ecosystem Exists for Verifiable Claims".

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-1
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-1
http://w3c.github.io/vctf/meetings/2016-01-28/#topic-1
http://w3c.github.io/vctf/meetings/2016-01-29-2/#topic-2
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-1
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0048.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0049.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0051.html

6.2 Current Technologies Are Not Readily Solving the Problem

A broad range of technologies that could be applicable to the Verifiable Claims Problem Statement were discussed during the interviews. Some of these technologies included OpenID 1.0, OpenID Connect, SAML, OAuth, Identity Credentials, WebDHT, Blockchain-like technologies, XRI/XDI, Hierarchical Deterministic (HD) Keys, and the JOSE stack. While it was identified that a few of these technologies could be a part of the solution to the Problem Statement, none of them were being employed at scale to effectively solve the problem today.

References:
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-2
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-4
http://w3c.github.io/vctf/meetings/2016-01-28/#topic-3
http://w3c.github.io/vctf/meetings/2016-01-29-2/#topic-6
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-3
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-5
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0048.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0049.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0051.html

6.3 W3C and IETF have an Important Role to Play

Each interviewee was asked which standards-setting organization should play a role in addressing the Problem Statement. Many of the existing standards organizations, like W3, IETF, OASIS, ITU, and others, were discussed. The suggestion that most of the interviewees made, unsurprisingly, was for the higher-level application layer work to be done at W3C and the lower-level protocol work to be done at IETF.

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-5
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-6
http://w3c.github.io/vctf/meetings/2016-01-28/#topic-5
http://w3c.github.io/vctf/meetings/2016-01-29-2/#topic-8
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-6

6.4 Clear Use Cases Exist for Verifiable Claims (including Payments)

Throughout the last four years, the Web Payments Community Group, the Credentials Community Group, and now the Verifiable Claims Task Force have been gathering use cases from a variety of different industries, including payments, financial, insurance, healthcare, government, and education. The use cases now include feedback from 51 different people and organizations (listed toward the front of this document). The current status of this work can be found here:

http://opencreds.org/specs/source/use-cases/

The raw use case input can be found here:

http://opencreds.org/presentations/2015/w3c-tpac/anonymized.html

and in the interviews:

https://docs.google.com/document/d/1dYup3KC2nak3LVTzyapr996TKxDj1w5Eyp4g13rQQBA/edit#heading=h.474cckqta3ic

6.5 Minimum First Step is to Establish a Way to Express Verifiable Claims

Many of the interviewers suggested that having a data model and syntax for the expression of verifiable claims as only part of the solution. Some of the interviewers asserted that the technology already exists to do this and that W3C should focus on vocabulary development. Others asserted that vocabulary development is already happening in focused communities (such as the Badge Alliance, the Credentials Transparency Initiative). Many of the interviewers suggested that the desirable outcome of standardization work is not only a data model and syntax for the expression of verifiable claims, but a protocol for the issuing, storage, and retrieval of those claims, but acknowledged that it may be difficult to convince W3C member companies to undertake all of that work in a single Working Group charter.

In the end, consensus around the question what kind of W3C charter would garner the most support seemed to settle on the creation of a data model and one or more expression syntaxes for verifiable claims. It was also suggested that another output of the Working Group should be a NOTE related to how the data model and syntaxes should be utilized in current protocols or if a new set of protocols would be required.

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-5
http://w3c.github.io/vctf/meetings/2016-01-28/#topic-4
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-5
https://lists.w3.org/Archives/Public/public-credentials/2016Feb/0012.html
http://w3c.github.io/vctf/meetings/2016-02-12/#topic-3
https://lists.w3.org/Archives/Public/public-credentials/2016Feb/0010.html
http://w3c.github.io/vctf/meetings/2016-01-29-2/#topic-7
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-4
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0048.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0049.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0051.html

7. Areas of Concern

7.1 Ensuring a Larger Vision is Effectively Communicated

Several interviewees raised a concern that if a limited Verifiable Claims Working Group Charter covering only data model and syntax was proposed that the larger vision of an interoperable verifiable claims ecosystem supporting issuing, storage, and retrieval would not be seen by the W3C Membership. If the full vision is not visible to W3C members, there was a concern that positive votes on an initial charter could be negatively affected. It was also noted that the ecosystem would not reach its full potential unless there was a protocol, with a browser API, for issuing, storage, and retrieval in addition to the verifiable claims data model and syntax.

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-5

http://w3c.github.io/vctf/meetings/2016-01-27/#topic-5
http://w3c.github.io/vctf/meetings/2016-02-12/#topic-3
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-4

7.2 Scalability of Trust Model with Thousands of Issuers and Consumers

There was a concern raised about decoupling issuers from identity providers resulting in tens to hundreds of thousands of different issuers and consumers. The current trust model that the Web uses, which is largely the Certificate Authority system, only has a few hundred issuers. There was concern that having hundreds of thousands of issuers would not be scalable. A counter-point was made that you can scale because there will be "islands of trust" in the verifiable claims ecosystem and not a "fully connected graph of trust". If the trust model is not found to scale, the theory is that it will scale back to the delegated trust model we have with the Certificate Authority system today with the added benefit that the system will still be user-centric and privacy enhancing.

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-5
http://w3c.github.io/vctf/meetings/2016-01-28/#topic-2
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-4
http://w3c.github.io/vctf/meetings/2016-02-12/#topic-1
https://lists.w3.org/Archives/Public/public-credentials/2016Feb/0010.html
http://w3c.github.io/vctf/meetings/2016-01-29-2/#topic-5
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0048.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0049.html

7.3 Compelling Business Models and Deployment Economics

There was a concern raised around the economics for reaching billions of users with a verifiable claims ecosystem. Verifiable claim issuers, holders, storage providers, and consumers each need to have a strong economic incentive for participating in the ecosystem. It should be noted that the concerns were raised primarily by people outside of the verifiable claim issuing, storage, and consumption industry.

Interviews and surveys have been done in that industry and it has been demonstrated that there is clear economic incentive to participate. The primary identified concern has been around a catch-22 related to issuing and consumption of verifiable claims. Organizations won't issue verifiable claims if they aren't being consumed and organizations that consume verifiable claims won't start consuming claims if they are not being issued. However, there are a number of large organizations participating in the Credentials Community Group that can kick-start the market, and have asserted that they will do so, but it is yet unknown if kickstarting the market will be enough to build a strong economic incentive feedback loop.

References:

http://w3c.github.io/vctf/meetings/2016-01-08/#topic-5
https://lists.w3.org/Archives/Public/public-credentials/2016Feb/0012.html
http://w3c.github.io/vctf/meetings/2016-02-12/#topic-1

7.4 Business Model for Basic Ecosystem Infrastructure

Several concerns have been raised around the need for new infrastructure technologies, such as self-sovereign/decentralized identifiers, website-based browser polyfills, and other core technologies that are needed to achieve the verifiable claims ecosystem protocol requirements (if a new protocol is needed). If core infrastructure is needed, which set of organizations are going to provide the funding to setup and run that core infrastructure? There is a group that is currently forming a non-profit specifically designed for this purpose and is already raising funding, but it is still unknown exactly how much funding will be required to bootstrap and run the basic infrastructure until a strong economic incentive feedback loop is created.

References:

http://w3c.github.io/vctf/meetings/2016-01-08/#topic-5
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-4
https://lists.w3.org/Archives/Public/public-credentials/2016Feb/0012.html
http://w3c.github.io/vctf/meetings/2016-02-12/#topic-5
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-1

7.5 Slow Evolution of Agent-Centric Designs

Agent-centric designs, such as SSL/TLS, web browsers and servers, and OpenID clients are slow to evolve and upgrade. Contrast this with super-providers that can force their clients to upgrade or risk being rejected. It is important to note that the user-centric, privacy-enhancing approach is certainly an agent-centric approach and thus carries with it the risks of slow evolution and upgrade cycles. Any group attempting the technical work should be aware of these risks and realities when designing the technology.

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-4

7.6 Long-Lived Identifiers, Key Management, and Revocation of Credentials

A few interviewees noted that they believed that long-lived identifiers are an anti-pattern and should be avoided due to negative privacy implications. Other interviewees noted that once key pieces of data are known about an individual (such as an email address), they are as effective as a long-lived identifier. Another issue with long-lived identifiers are the possibility of long-lived verifiable claims associated with these identifiers and the revocation rules around those verifiable claims.
A few others noted that key management is a particularly catastrophic user requirement as well. These items have been discussed at length in the Credentials Community Group over the past 18 months and there are approaches to addressing each one of these concerns. However, specific decisions can’t be made unless there is a Working Group that will decide on the proper balance of techniques and technologies that will address the Problem Statement, use cases, and identified requirements.

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-3
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-5
http://w3c.github.io/vctf/meetings/2016-02-12/#topic-2
https://lists.w3.org/Archives/Public/public-credentials/2016Feb/0010.html
http://w3c.github.io/vctf/meetings/2016-01-29-2/#topic-5
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0048.html
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0049.html

7.7 Protection Against Fraud and Abuse

It was noted that any work in the area should understand the fraud and abuse models for stolen/unlocked devices that have access to one’s verifiable claims. If a verifiable claim is used to do damage, who is liable? The issuer, holder, storage provider, or organization consuming the verifiable claim? It was suggested that legal input should be requested on this particular topic.

References:
http://w3c.github.io/vctf/meetings/2016-01-08/#topic-2

7.8 Re-using Previous Work Efforts

It has been suggested that previous work efforts in the area, such as OAuth, the JOSE stack, SAML, OpenID Connect, JSON, and a variety of other technologies should be used to address the problem statement, use cases, and requirements. There is consensus that where a technology exists that can be re-used, it should be re-used. However, it has also been noted that a deeper technical analysis on which technologies should be used is the purview of a Verifiable Claims Working Group, if one were to be created.

References:

http://w3c.github.io/vctf/meetings/2016-01-28/#topic-3
http://w3c.github.io/vctf/meetings/2016-01-27/#topic-2
http://w3c.github.io/vctf/meetings/2016-01-29-2/#topic-6
http://w3c.github.io/vctf/meetings/2016-01-29-1/#topic-3
https://lists.w3.org/Archives/Public/public-webpayments-ig/2016Jan/0049.html