Smart Consent Protocol

A White Paper from Rebooting the Web of Trust III Design Workshop

By Dr. Shaun Conway, Lohan Spies, Jonathan Endersby, Tim Daubenschütz

Consent https://consent.global @globalconsent **COALA IP** https://coala.global

PERSONAL DATA AS DIGITAL INTELLECTUAL PROPERTY

Personal Data are valuable resources for creating digital intellectual property (IP). Rights over this IP have generally been unclear, resulting in systematic abuse or unfair use of people's personal data by third parties. But new regulations are changing this - most notably, the European Union General Data Protection Regulation (EU GDPR). Third parties must now obtain explicit and documented consent from people (data subjects) to collect, process, store or disclose their personal data. A specification for operationalising these regulatory requirements, using digital Consent Receipts, is being developed through the Consent and Information-Sharing Working Group of the Kantara Initiative. In a parallel effort, COALA-IP has developed a blockchain-ready, community-driven generic protocol for intellectual property licensing that applies the Linked Content Coalition (LCC) framework to "unify digital rights data management". This paper proposes a decentralised Smart Consent protocol for managing personal data as intellectual property that combines elements of the COALA-IP Specification for Digital Intellectual Property with the specification for Digital Consent Receipts.

RIGHTS AND WRONGS

The World Economic Forum reported in 2011 that 'Personal Data is becoming a new economic asset class, a valuable resource for the 21st century that will touch all aspects of society'. This would logically include personal data that are used to assert verifiable claims relating to digital identity. Like the

Microsoft
Blockstream
Protocol Labs
TIERION
NETKI
Sponsors for the Rebooting the Web of Trust III Design Workshop

many other derivatives of personal data, verifiable claims can increase in value the more they are relied on by third parties and as they gain provenance through attestations and use over time. Innovations such as Artificial Intelligence and Behavioural Analytics will further enrich these personal digital assets with additional intellectual properties and embedded value that should benefit individuals, organisations and society.

However, it is likely that corporations will continue to benefit most from the commercial uses of Personal Data, as these algorithms and data processing capabilities are mostly proprietary and opaque. This especially biases toward entities that have already amassed personal data with few IP controls and that have developed the most efficient methods of extracting people's personal data, without real consent. This means that entities acting out of economic or political self-interest have become the primary aggregators and processors of Personal Data. Decisions about how to use or share the data are therefore based on the likelihood of profitability and return on investment.

It seems wrong that market dynamics should determine the types, quality and availability of Personal Data and who gets the privileges of benefiting from this. This does not primarily serve the interests of the person who is the original source or subject of the data, nor the common good of society.

A significant part of the problem is lack of effective control over who owns the rights to this valuable Personal Data and how it gets used. This often has negative (even if unintended) consequences. For instance, people risk being economically exploited or socially disadvantaged by companies using their personal data for profiling and filtering. How businesses are currently using our personal data is creating enormous trust deficits in society, which has broader economic consequences.

Without addressing rights of ownership and use of personal data, the potential for it to be a valuable resource will continue to be compromised. But it seems implausible that this situation will be resolved by central authorities or regulatory controls.

THE CASE OF BENJAMIN

Benjamin was born with the hereditary disorder Duchenne Muscular Dystrophy. His physician, suspecting the diagnosis, orders a DNA test from a commercial clinical laboratory. A digital file of unique genetic code sequences is generated, using a patented genetic probe. The resultant data are processed by a proprietary algorithm matching Benjamin's data fingerprint against a reference database extracted from clinical trial participants who were paid for their data.

The result of this investigation is a positive diagnosis. Benjamin is now identified as "A person living with Duchennes". This claim is verified by a licensed geneticist who provides an electronically signed report back to Benjamin's physician.

The various digital assets that have been created through the course of this investigation have direct value to a number of prospective users of Benjamin's personal data. For instance, health insurers will pay out healthcare claims on the basis of Benjamin's having been verified as 'A person living with Duchennes'. Copies of Benjamin's Personal Data and various data elements of his verifiable claim persist in multiple third-party systems. These are mostly regulated by health information, privacy and other protections.

Now a pharmaceutical company that is developing a candidate CRISPR gene therapy wishes to use Benjamin's data for its research. How should consent be obtained to access the various datasets and their derivative products? Who should legitimately give this consent? And what would be fair remuneration for the new rights of use to compensate the various parties that have added valuable intellectual property to *Benjamin's* digital assets, as they have passed through the value chain?

Unresolvable assumptions and uncertainties about the ownership and rights of use of Benjamin's data are likely to diminish its potential commercial, academic and personal value.

There are many incongruent and deficient laws, regulations and technologies relevant to how personal data gets processed and used. This does not currently serve the interests of the people generating or using Personal Data. As the amount and scope of digitized personal information grows, it will become even more difficult for multiple entities to responsibly hold significant parts of this information and comply with personal data protection regulations, without having a persistent record of their rights to use, process, hold and store each personal data asset.

The most logical solution would be for a form of digital certificate to be encoded with each personal data asset, referenced against a public record.

A RIGHTS WAY FORWARD

COALA IP (Coalition Of Automated Legal Applications, Intellectual Property) was formed to design and implement a free and open specification for handling digital licensing of intellectual property. Its goals are to establish open, free, and easy ways to claim attribution, add metadata, license works, mediate IP disputes, and authenticate claims of others. The group believes that there should be global agreement at the data level without the need for centralised control. COALA IP extends the LCC Framework to represent IP Rights digitally with a standardised data model. For instance, this uses the RDF (Resource Description Framework) standard to record assertions in a JSON-LD format.

Consent Receipts are part of a broader Consent Framework for individuals to track and control the use of their personal information by third parties. This is being developed as a set of open-source specifications through the Consent and Information-sharing Working Group (CISWG) of the Kantara Initiative. The group's vision for Open Consent is that: *Once transparency over data control is achieved and people are able to manage consent holistically, there will be more control and trust in the way people share information and trust, enabling people to explicitly assert preferences, attributes, and manage pseudonymity from a trusted notice, consent and privacy framework.*

A Consent Receipt records a standard set of legal, social and contextual parameters relating to an information-sharing transaction. Standardisation of the record should promote consistent consent practices, interoperability between systems and services, and a globally accepted proof of consent.

The technical specification for Consent Receipts (still under development) defines the minimum required information elements for the electronic version of a Consent Receipt, with a schema for encoding this in a standard JSON data structure.

By combining elements of these specifications, when personal data are shared with a record of consent, the terms of use of these data can be embedded in the Consent Receipt. This record can stored with an immutable proof (hash value) that cannot be repudiated, on a decentralised public ledger. This becomes a powerful reference to the rights of ownership and use of these digital assets.

The advantage of expressing Consent Receipts in an RDF schema is that it allows the semantic mapping of related but domain-specific concepts. In our case Intellectual Property and Personal Data:

As an example...

COALA IP describes a digital right in this JSON-LD format:

```
{   "@type": { "/": "<hash pointing to RDF-Schema of Right>" },
    "usages": "all|copy|play|stream|...",
    "territory": { "/": "<hash pointing to the Place>" },
    "context": "inflight|inpublic|commercialuse…",
    "exclusive": "true|false",
    "manifestation": { "/": "<hash pointing to the Manifestation>" },
    "license": { "/": "<hash pointing to the License>" }
}
```

Consent Receipts describe an agreement for the use of Personal Data in the JSON format:

```
{
    "jurisdiction": { "/": "<hash pointing to a Place>" },
    "iat": "2007-12-24T18:21Z",
    "moc": "tbd",
    "dataController": {
        "onBehalf": "true | false",
        "controller": {
```

```
            "contact": { "/": "<hash pointing to a ContactPoint>" },
            "address": { "/": "<hash pointing to a PostalAddress>" },
            "email": "email@example.com",
            "phone": "+123456789",
    },
    "policyUri": { "/": "<hash pointing to a Policy>" },
    "services": {
        "serviceName": "Transactional Banking",
        "purposes": {
            "consentType": "Required | Explicit, Opt-in | ...",
            "purposeCategory": "tbd: commercial, research, compliance, ...",
            "piiCategory": ["Marketing", "Personalized Experience", "..."],
            "nonCorePurpose": "true | false",
            "thirdPartyDisclosure": "true | false",
            "thirdPartyName": "Example Inc.",
        }
    },
    "sensitive": "true | false",
    "sub": { "/": "<hash pointing to a Person or Organization>" },
    "spiCat": "religious believe | criminal convictions | ..."
}
```

Both the COALA IP Rights model and Consent Receipt specification include references to documents that are intended to be read by humans, such as: `license` pointing to a document that describes terms and conditions for using and distributing a digital work, or `policyUri` that points to the legal description of an organization's privacy policy.

Additional similarities include:

COALA IP Right	*Consent Receipt*	*Comment*
territory	jurisdiction	The geographical agreement is valid.
usage & context	purpose	Context in which the agreement is valid and purposes for which the information will be used.

COALA IP assumes that license data will be recorded in transactions on immutable ledgers. Using this method for recording Smart Consent makes it unnecessary to declare a `jti` timestamp and `publicKey`as part of the Consent Receipt.

The Consent Receipt can be extended to include a digital right, related to a specific purpose. For instance, if the primary purpose is to make a clinical diagnosis and the non-core purpose is to use the personal data for clinical research (no charge for academic research, compensation for commercial pharmaceutical research), the Consent Receipt could include:

```
javascript {
  "@type": {
    "/": "<hash pointing to RDF-Schema of Right>"
  },
  "usages": "clinical research",
  "territory": {
    "/": "<hash pointing to a Place>"
  },
  "context": "academic|commercial",
  "exclusive": "true|false", //  ...
  "manifestation": {
    "/": "<hash pointing to the Manifestation>"
  },
  "license": {
    "/": "<hash pointing to the License>"
 }
}
```

TECHNICAL SPECIFICATIONS

Specifications for various rights assignments will need to be developed or adapted from existing frameworks. These should be flexible. For instance, to enable temporary transfer of the rights of use for a specific purpose (Benjamin grants a right of use over his personal data to the pharmaceutical company for the purpose of clinical research) or in another context, to extend ownership over a set of Personal Data assets (Benjamin agrees that his data can also be owned by the company that provides the diagnostic algorithm, so they can continue to enhance and sell their database).

Ideally, this approach should support a model that enables private and commercial creators or processors of Personal Data to make the resultant digital assets more freely available to others. This is not about restricting usage of Personal Data by imposing a system of Digital Rights Management! A standardised technical specification for recording Intellectual Property Rights over Personal Data should promote new economic models for sharing and generating benefits from personal data, whilst protecting individual privacy. For instance, services that further enrich personal data or create premium derivative assets from this (including verified claims), could be fairly and transparently compensated. This should increase trust and promote growth in the personal dataeconomy.

UNRESOLVED ISSUES

Currently, Consent Receipts are designed to be mutable models (e.g.`purposeTermination`, defining the revocation of a consent). This makes a transformation in some cases difficult. However, concepts of revocation on immutable ledgers are under development by various influential organizations (e.g. W3C, WOT, ...).

Further work needs to be done in standardizing the Consent Notice Receipt Specification vocabulary to achieve greater alignment.

NEXT STEPS

This paper was drafted after brief discussions with Trent McConaghy (COALA IP Working Group) and Mark Lizar (Consent and Information Sharing Working Group, Kantara Initiative) and by reviewing the draft specifications from each group. It is intended to promote discussion amongst stakeholders across these and other groups about whether there is a real need for additional technical specifications to specifically address IP rights of Personal Data.

This could initiate a collaboration to develop and test new standards through a process that defines a minimum viable schema definition to extend the RDF Schema definitions of COALA IP and embed these as additional elements in the specification for Digital Consent Receipts.

REFERENCES

COALA IP Specification, https://github.com/coalaip/specs

How Blockchains can support, complement, or supplement Intellectual Property. Working Group on Intellectual Property, COALA IP (May 2016), https://docs.google.com/viewer?url=http%3A%2F%2Fcoala.global%2Fwp-content%2Fuploads%2F2016%2F06%2FCOALA-IP-Report-May-2016.pdf

Consent Notice Receipt Specfication (Version 0.9.1). Kantara Initiative, Consent and Information Sharing Working Group, 2 October 2016, https://docs.google.com/document/d/1-n06avXzwdYM6SeF1siUwNyhQIONreYMglXm9CZ7J0c/edit

Personal Data: The emergence of a new asset class. World Economic forum, 2011, https://docs.google.com/viewer?url=http%3A%2F%2Fwww3.weforum.org%2Fdocs%2FWEF_ITTC_PersonalDataNewAsset_Report_2011.pdf

Additional Credits

Authors: Dr. Shaun Conway, Lohan Spies, Jonathan Endersby, Tim Daubenschütz

About Rebooting the Web of Trust

This paper was produced as part of the **Rebooting the Web of Trust III** design workshop. On October 19th through October 21st, 2016, over 40 tech visionaries came together in San Francisco, California to talk about the future of decentralized trust on the internet with the goal of writing 3-5 white papers and specs. This is one of them.

Workshop Sponsors: Blockstack, Microsoft, Netki, Protocol Labs, Tierion

Workshop Producer: Christopher Allen

Workshop Facilitators: Christopher Allen and Brian Weller, additional paper editorial & layout by Shannon Appelcline, and additional support by Kiara Robles and Marta Piekarska.

What's Next?

The design workshop and this paper are just starting points for Rebooting the Web of Trust. If you have any comments, thoughts, or expansions on this paper, please post them to our GitHub issues page:

https://github.com/WebOfTrustInfo/rebooting-the-web-of-trust-fall2016/issues

The next Rebooting the Web of Trust design workshop is scheduled for Spring 2017 in Paris, France. If you'd like to be involved or would like to help sponsor these events, email:

ChristopherA@LifeWithAlacrity.com