INTERNATIONAL COMPUTATIONAL LAW JOURNAL

CODEX SOURCE ARCHIVE

The Data Rights Protocol: Threading Privacy Rights into the Internet

Sebastian Zimmeck

Computational Law Report · Release 4.0 · 2022-01-18

Reader edition prepared by The Research University under CC BY 4.0. This cover was added; the original PDF pages follow unchanged. Original authorship, publication identity and source links are preserved. This independent republication is not sponsored or endorsed by Stanford University, Stanford Law School, CodeX or MIT.

Volume 1 · Issue 6 · Rights, Contracts and Digital Institutions

Original source and license evidence are recorded on the article landing page.

MIT Computational Law Report

The Digital Rights Protocol: Threading Privacy Rights into the Internet

Sebastian Zimmeck

Published on: Jan 18, 2022

License: Creative Commons Attribution 4.0 International License (CC-BY 4.0)

ABSTRACT

The Data Right Protocol (DRP)[1] has the goal of simplifying privacy rights requests for everyone in the Internet ecosystem by providing standardized API endpoints. While new privacy laws, such as the California Consumer Privacy Act (CCPA), give users privacy rights to opt out, access, or delete their data, those rights are difficult to exercise in practice. Burdens do not only exist for users but for companies as well. The processing of rights requests is often manual and, thus, time-consuming and costly.[2] In short, privacy rights are increasingly provided, but they are difficult to exercise and manage for everyone involved in the request process. One of the main reasons is the lack of usable privacy rights technologies. There is no uniform way of how privacy rights ought to be requested and fulfilled on the Internet. DRP intends to simplify and streamline this process.

Privacy Rights on the Internet

DRP's main idea is to thread privacy rights into the structure of the Internet. Ideally, any system would have privacy built in right from the start, but since the Internet and its data-centric business models evolved with only few guard rails for privacy, we are now at a point where we have to retroactively construct and fit those into its already existing structure. The good news is that there is considerable momentum for doing so. Privacy law-making activity, recognition by companies that good privacy practices make good business sense, and an increase in public awareness of privacy issues contribute to this development. Since the General Data Protection Regulation (GDPR) became effective in 2018 and the success of the CCPA ballot proposition in the same year, privacy lawmaking has gathered steam both internationally as well as stateside. Germany passed the new Telecommunications and Telemedia Data Protection Act.[3] In the US, in addition to California's CCPA, Colorado and Virginia passed privacy acts as well.[4] The CCPA gives consumers three rights that are also recognized in Colorado and Virginia:

- Do not sell and do not share: the right to opt out of the sale and sharing of personal information
- Access: the right to know what personal information a company has collected
- Deletion: the right to have personal information deleted

This set of rights is also initially supported by DRP. Other laws provide more extensive rights. For example, the GDPR also provides the right to data portability. In order to

accommodate this right as well as other rights DRP is being drafted with room for extensibility.

Challenges for Establishing an Internet Privacy Rights Infrastructure

Bringing privacy rights technologies into the Internet infrastructure requires addressing various hard challenges:

- Usability: An overarching challenge that touches upon essentially every facet of making it easy for users to exercise their privacy rights is usability. For example, giving users the option to select their rights on a per-site basis without appropriate default settings or options to generalize their privacy decisions for sets of sites will be a serious usability problem as the cookie banner experience shows. DRP aims for an abstract design that is flexible enough to accommodate such usability features.
- Authentication: To a varying degree all rights require some form of authentication of the requesting user; the right to access and delete to a greater extent than the right to opt out. Otherwise, personal information could be mistakenly disclosed to the wrong users. Indeed, the lack of identity verification and fraud protection are some of the main reasons why companies are denying rights requests.[5] The envisioned reuse of OpenID Connect for DRP purposes is one way of addressing these authentication challenges.
- Different jurisdictions: Determining the geographic location from which a request originates is critically important as different jurisdictions often require a different handling of rights or even have a different set of rights altogether. Some companies may decide to give users more rights than they would have in their jurisdiction. However, in principle, users' geographic locations must be known to identify which set of rights is applicable to them. Thus, DRP's Data Rights Exercise endpoint is designed with a "regime" field that can have values like "ccpa" or "gdpr."
- Complexity of communication: The communication between a user and a company must account for what should happen in case of incomplete, canceled, non-matching, or erroneous requests as well as requests for which no data is available. DRP is being designed with a response mechanism for answering rights requests as necessary.
- Conflicting rights: There may be requests that would have to be denied because of other users' rights or mandatory laws. For example, some data may need to be retained and could not be deleted if it is related to the investigation of a crime. Here, DRP provides various states and reasons to notify the requesting user.

These are just some of the challenges and how DRP addresses those. The challenges are substantial and require broad participation of all stakeholders in the process of advancing DRP towards a full Internet Privacy Infrastructure.

Standardizing and Implementing DRP

As the law is catching up with the reality of data-centric business models, we are now at a point at which many users are provided with rights to control their personal information. However, there is a gap between the law and the implementation of technologies that would make privacy rights effectively usable for users. Privacy standards at standards bodies such as the W3C or IETF are the right venue to take the law and translate it into implementable technologies. The CCPA regulations explicitly provide room for such standards[6] and the same is true for the German Telecommunications and Telemedia Data Protection Act.[3] In order to make progress, law, technology, and business models have to go hand in hand. However, legislators and regulators are not in a position to go down the path of technological implementation decisions. That is the job of the standard communities. They have to take it from there and fill the law with life. Now is the time to infuse the Internet with privacy standards.

Standardization of privacy rights is necessary but not sufficient. It is necessary for the same reasons as standards in other areas are necessary: a uniform approach makes sure that everyone can communicate with one another. With a standard it can be ensured that every site that follows it can process privacy rights requests and fulfill them across all endpoints connected in the ecosystem. We need an approach that allows implementation across sites and for all companies that collect data on the Internet. Exercising privacy rights site-by-site does not work.[7] Users should be able to opt out, access, or delete their data from multiple or all sites if they wish. Standardizing privacy rights is necessary to make exercising privacy rights usable. However, a standard is of no use if it is not applied in practice. A standard must be lived. In addition to putting it into words, the standard must be implemented. Thus, DRP follows a strategy of implementation and standardization in parallel involving publishers, consent management platforms, and other stakeholders in the Internet ecosystem.

DRP is intended to be easy to understand for users and easy to implement for the different types of implementers. One way in which DRP aims to foster adoption is by reuse of existing technologies. For example, to address the challenge of authenticating

users the authentication layer of OAuth 2.0 — OpenID Connect[8] — may be used as a familiar way for implementers to authenticate users. The already existing experience of implementers with OpenID Connect or other widely deployed technologies lowers the barrier of implementation. This principle holds in general. Whenever possible, it is a good idea to reuse existing technologies that are already widely adopted if their purpose can be leveraged in the context of DRP. To be sure, they may not always completely solve the challenge at hand. For example, in the case of authentication, not all sites may actually authenticate their users via a login but rather may just use a cookie ID. In those cases OpenID Connect may not be the right solution. The reason is that users should be authenticated by a site for purposes of rights requests by whatever means the site usually uses to authenticate their users. Information collection from users should be minimized, and they should never have to input information that they have not already provided [9].

DRP's Relationship to Global Privacy Control (GPC)

Various privacy rights preference technologies are currently in development.[10] As those technologies mature they should be kept interoperable as far as possible. As one of the most closely related technologies, GPC[11] shares various similarities with DRP. Both technologies are complementary to each other. Their common idea is the automation of privacy rights processing. In terms of their practical approach, they also both kickstart adoption with an initial group of implementers. Both cover CCPA Do Not Sell and Do Not Share rights. Thus, users can use either DRP or GPC to opt out of the sale or sharing of personal information per the CCPA. Both are intended to be extensible as to the rights they cover. However, GPC is emphasizing the opt out right while DRP takes a broader approach covering access and deletion rights as well. GPC uses a simple binary signal — either a user is opted out or not — that a browser or extension vendor can easily add to existing privacy settings without substantially increasing the fingerprinting risk or causing lots of sites to break.

DRP and GPC are approaching the automation of privacy rights from slightly different angles: DRP is a business-facing standardization approach aimed at Privacy Infrastructure Providers (PIPs); GPC is more geared towards user agent implementation in browsers and browser extensions. DRP is broader than GPC with its encoding of a set of standardized request and response data flows. It enables users to exercise rights provided under a broad set of laws and regulations, including the CCPA, GDPR, and voluntary bases, and receive affirmative responses in standardized formats. DRP is aimed to be a standard integrable with the emerging ecosystem of

data rights middlewares, agent services, and automation tool kits. A core element of DRP is its focus on PIPs. For example, many websites integrate consent management platforms (CMPs) to show their users consent banners and store consent choices. DRP provides CMPs and other PIPs with a protocol to process user rights requests in a standardized way.

Implementers can adopt GPC without adopting DRP and vice versa. Adoption of DRP is entirely voluntary while GPC adoption is mandatory for many California businesses. Per the California Attorney General, GPC signals are binding opt out requests.[12] GPC and DRP follow their own release timelines independently. To avoid blockage they are not dependent on one another. Various use cases illustrate the complementary nature of DRP and GPC:

- To the extent websites are respecting a user's GPC signals, there is generally no need for the user to exercise access or deletion rights via DRP.
- After a user has sent an opt out signal via GPC, they can use DRP to verify that the personal information to which the opt out applies is not being sold to or shared with other companies. Verifying whether an opt out really occured can also help with the enforcement of privacy rights. For example, a user who submitted an order form with GPC turned on and later finds the information they entered via a DRP data access request at another company, they can easily report a CCPA violation,[13] if any.
- DRP can help a user to claim their access and deletion rights in situations where Do Not Sell or Do Not Share requests cannot be sent. For example, a retailer may have collected data from a user in a physical retail store and passed it on to a data broker. The user can now request access and deletion of the data from the data broker.

The Future of Privacy Rights on the Internet Starts ... Now

Over the next months and years we will likely see more legislative and regulatory activity — both stateside and internationally — to provide privacy rights to Internet users. This legal activity is the basis for making the Internet more private, which is not only inherently important from a human rights perspective but also a prerequisite for making the Internet more safe. However, the law is not enough. With just the law in place, the work is far from being done. Without usable privacy rights implementations it is not possible for users to effectively claim their privacy. This is the job of DRP. With the parallel strategy of implementing and standardizing privacy rights into usable technologies, DRP offers the promise of bringing the law to life. We have the chance to make the Internet more private. The time is now. Let's get to work!

> *The author thanks Dazza Greenwood, Don Marti, Boaz Sender, Ryan Rix, Justin Brookman, and Ginny Fahs for comments on the draft of this blog post.*

Sebastian Zimmeck is an assistant professor of computer science at Wesleyan University's Department of Mathematics and Computer Science. He is working on web and mobile app privacy and is directing the privacy-tech-lab. Sebastian is one of the creators of Global Privacy Control and is on the Technical Advisors Board of the Data Rights Protocol.

Citations

[1] Consumer Reports, Data Rights Protocol.

[2] Gartner estimates that manually processing a single data rights request costs $1,406. Cited from DataGrail, The State of CCPA.

[3] German Telecommunications and Telemedia Data Protection Act. See, in particular, Section 26.

[4] See the Colorado Privacy Act and Virginia's Consumer Data Protection Act.

[5] privacy-tech-lab CCPA metrics.

[6] The newly established California Privacy Protection Agency is tasked with adopting regulations to define the requirements and technical specifications for an opt out preference signal per California Civil Code, Sections 1798.185(a)(19)(A) and 1798.185(d).

[7] Maureen Mahoney, Consumer Reports, California Consumer Privacy Act: Are Consumer's Digital rights protected?

[8] OpenID Connect.

[9] Sean O'Connor, Ryan Nurwono, Aden Siebel, Eleanor Birrell, (Un)clear and (In)conspicuous: The Right to Opt-out of Sale under CCPA.

[10] Other technologies beyond DRP and GPC include the Advanced Data Protection Control and industry efforts, such as the Interactive Advertising Bureau's (IAB) US Privacy String or the IAB Europe's Transparency Consent Framework.

[11] Global Privacy Control.

[12] State of California Department of Justice, CCPA Enforcement Case Examples ("The business neither imposed a service provider contractual relationship on these third parties, nor processed consumers' requests to opt-out that were submitted via a user-enabled global privacy control, e.g., a browser extension that signaled the GPC. After being notified of alleged noncompliance, [...]").

[13] State of California Department of Justice, Consumer Privacy Interactive Tool.