NIST Internal Report
NIST IR 8588 ipd

A Community-Driven Differential Privacy Deployment Registry

Initial Public Draft

Micah Altman
Sharon Ayalde
Rachel Cummings
Damien Desfontaines
Jack Fitzsimons
Elena Ghazi
Andrew Gruen
James Honaker
Gary Howarth
Nitin Kohli
Chuck McCallum
Priyanka Nanayakkara
Joseph P. Near
Robert Pisarczyk
Salil Vadhan

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8588.ipd

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Internal Report
NIST IR 8588 ipd

A Community-Driven Differential Privacy Deployment Registry

Initial Public Draft

Gary Howarth
Privacy Engineering Program
Information Technology Laboratory, NIST

Micah Altman
Massachusetts Institute of Technology

Sharon Ayalde
Elena Ghazi
Chuck McCallum
Priyanka Nanayakkara
Salil Vadhan
Harvard University

Rachel Cummings
Columbia University

Damien Desfontaines
Hiding Nemo

Jack Fitzsimons
Robert Pisarczyk
Oblivious

Andrew Gruen
Working Paper

James Honaker
Mozilla Co.

Nitin Kohli
University of California, Berkeley

Joseph P. Near
University of Vermont

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8588.ipd
September 2025

U.S. Department of Commerce
Howard Lutnick, Secretary

National Institute of Standards and Technology
Craig Burkhardt, Acting Under Secretary of Commerce for Standards and Technology and Acting NIST Director

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at https://csrc.nist.gov/publications.

NIST Technical Series Policies
Copyright, Use, and Licensing Statements
NIST Technical Series Publication Identifier Syntax

How to cite this NIST Technical Series Publication:
Altman M, Ayalde S, Cummings R, Desfontaines D, Fitzsimons J, Ghazi E, Gruen A, Honaker J, Howarth G, Kohli N, Mulligan D, Nanayakkara P, Near JP, Pisarczyk R, Vadhan S (2025) A Community-Driven Differential Privacy Deployment Registry. (National Institute of Standards and Technology, Gaithersburg, MD), NIST IR 8588 ipd. https://doi.org/10.6028/NIST.IR.8588.ipd

Author ORCID
Micah Altman: 0000-0001-7382-6960
Sharon Ayalde: 0009-0004-7527-2809
Rachel Cummings: 0000-0002-1196-1515
Damien Desfontaines: 0000-0002-9664-0963
Jack Fitzsimons: 0000-0001-6761-5315
Elena Ghazi: 0009-0000-5470-8379
Andrew Gruen: 0009-0006-6516-9730
James Honaker: 0000-0002-5404-539X
Gary Howarth: 0000-0002-3587-0546
Nitin Kohli: 0009-0008-9330-2978
Priyanka Nanayakkara: 0000-0002-0597-6657
Joseph P. Near: 0000-0002-3203-3742
Robert Pisarczyk: 0000-0001-5059-1599
Salil Vadhan: 0000-0002-4059-4072

Public Comment Period
September 17, 2025 – November 14, 2025

Additional Information
Additional information about this publication is available at https://csrc.nist.gov/pubs/ir/8588/ipd, including related content, potential updates, and document history.

Submit Comments
PETs@nist.gov

Abstract

There is a need in the community of privacy practitioners for a trustworthy, collaborative shared database of differentially private deployments, to help foster norms about best practices. We propose a shared governance resource to (1) help industry grow to consensus on best practices, (2) provide public snapshots of the privacy landscape so that regulators can judge new deployments in context and shape guidance accordingly, and (3) incentivize industry to make their choices public. We describe an initial schema to systematize this information and an editorial and governance process to ensure this information is reliable.

Keywords

Privacy; Differential Privacy; Privacy-Enhancing Technology; Data Governance

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems.

Call for Patent Claims

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this Information Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be directly stated in this ITL Publication or by reference to another publication. This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to this ITL draft publication and of any relevant unexpired U.S. or foreign patents.

ITL may require from the patent holder, or a party authorized to make assurances on its behalf, in written or electronic form, either:

1. assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or

2. assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this ITL draft publication either:

   (a) under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or

   (b) without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure that the commitments in the assurance are binding on the transferee, and that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

Such statements should be addressed to: PETs@nist.gov

Note to Reviewers

This draft proposes a structure for a new NIST resource, a voluntary differential privacy deployment registry. This registry seeks to serve as a communal body of knowledge to encourage the judicious application of differentially private implementations. The authors invite the differential privacy community to provide constructive feedback on maximizing the utility of such a resource. We particularly invite feedback on the following questions:

- Is the proposed schema responsive to foreseeable use cases, and, if not, what additions or changes are required?
- Are the mechanisms for reporting edits or updates sufficient and workable for practitioner needs?

Table of Contents

List of Tables

List of Figures

Author Contributions

Micah Altman: Writing- Original draft preparation, Data curation. **Sharon Ayalde**: Writing- Original draft preparation, Supervision. **Rachel Cummings**: Writing- Original draft preparation. **Damien Desfontaines**: Writing- Original draft preparation, Data curation. **Jack Fitzsimons**: Conceptualization, Funding acquisition, Writing- Original draft preparation, Data curation, Software, Supervision. **Elena Ghazi**: Writing- Original draft preparation, Data curation. **Andrew Gruen**: Writing- Original draft preparation. **James Honaker**: Conceptualization, Writing- Original draft preparation, Data curation, Software, Supervision. **Gary Howarth**: Conceptualization, Writing- Original draft preparation, Supervision. **Nitin Kohli**: Conceptualization, Writing- Original draft preparation. **Chuck McCallum**: Visualization, Data Curation. **Priyanka Nanayakkara**: Conceptualization, Methodology, Writing- Original draft preparation, Data curation, Software, Supervision. **Joseph P. Near**: Writing- Original draft preparation, Software. **Robert Pisarczyk**: Writing- Original draft preparation, Funding acquisition. **Salil Vadhan**: Conceptualization, Funding acquisition, Writing- Original draft preparation, Supervision.

1. Value of the registry

In 2019, Cynthia Dwork, Nitin Kohli and Deirdre Mulligan wrote "*Differential Privacy in Practice: Expose your Epsilons!*" [1], which argues that details about deployments of differential privacy should be collaboratively collected into a public resource database, they titled "the Epsilon Registry". From this, practitioners could share the choices they made for implementation—in particular the authors were interested in fostering common knowledge about the level of the privacy-loss parameter, epsilon, used in different settings. They wrote:

> *"…there is a need for shared learning amongst the differential privacy community. To serve these purposes, we propose the creation of the Epsilon Registry—a publicly available communal body of knowledge about differential privacy implementations that can be used by various stakeholders to drive the identification and adoption of judicious differentially private implementations."*

Since then, industry blogs and shared spreadsheets have begun collecting details, mostly scraped together from publicly available sources and shared notes [2]. These limited data points are becoming the de facto benchmarks against which engineers are designing differentially private systems.

There is growing need to realize the original goal of "Expose your Epsilons!", both to practitioners in making reasonable choices, and to foster and accelerate industry norms for privacy. Bringing together deployment data in a public and common format, with high-quality moderation will help speed industry consensus. A registry will ease adoption of differential privacy by establishing norms of use and application-specific guidance for parameter choices. A deployment registry will aid legal and policy reviewers to compare how new proposals sit within the landscape of other implementations in industry. The ability to compare with other deployments may also motivate systems to increase their privacy protections. Additionally, the focal nature of such a database, may provide industry more incentives and an increased norm to openness about their own deployments. To reflect the broader scope envisioned here, we adopt the name Differential Privacy Deployment Registry (DP Deployment Registry).

1.1. Value to Practitioners

Industry is in the early stages of adoption of differential privacy (DP) technologies. As such, privacy practitioners often meet challenges when trying to design and deploy formal privacy methods to a concrete use case. There are two kinds of such challenges: resistance from internal stakeholders, and open questions encountered in the course of the project.

The first kind of challenge is convincing stakeholders of the feasibility of using differential privacy: modern approaches to data privacy might still be perceived as experimental and unproven. Will differential privacy work for our specific data and query workload? Will

downstream data users get enough utility if we add noise to the computation results? Will sensitive data be adequately protected in a way that will be recognized by auditors and policy makers? The value of the DP Deployment Registry to a practitioner trying to answer these kinds of questions is obvious: if other organizations have successfully deployed this technology, it suggests that a similar approach is likely to also work for their use case. The more comprehensive and diverse the DP Deployment Registry grows over time, the more likely it is that potential practitioners will have similar examples to point to, and successfully make the case for investing in formal privacy approaches.

The second kind of challenge happens when trying to go from inception to deployment. A privacy practitioner will have to grapple with a number of difficult questions. How to quantify utility and determine whether the output data is fit for use? What methods will perform best to optimize the privacy-utility trade-off? What privacy parameters provide an adequate level of protection? How to complement a formal privacy analysis with an empirical evaluation of the potential leakage? What software libraries are available?

A DP Deployment Registry that includes details and rationale about the choices made by other practitioners will be a valuable source of information to guide and inform their own choices. It will also help new practitioners find other privacy experts that were involved in previous differential privacy projects, allowing them to ask questions and get feedback or guidance.

In both cases, the DP Deployment Registry catalyzes a virtuous cycle: helping practitioners roll out differential privacy to solve real-world use cases would then lead to successful deployments, and adding those deployments to the DP Deployment Registry would add to its value for future practitioners.

1.2. Value to Policy Makers

In a similar fashion, in our experience, policy makers face many of the same questions and encounter the same obstacles, from their mirrored perspective. Policy makers play a central role in shaping the norms and expectations around data protection. Their guidance documents and assessments often set the benchmarks that organizations must meet to demonstrate responsible data use. Yet these decisions are rarely made in a vacuum: to evaluate whether a deployment is sufficiently protective, policy makers must consider how it compares to the state of the art—what other organizations are doing, which practices are widely accepted, and where the consensus on adequacy lies. Without visibility into what has already been attempted and the reasoning behind the choices, policy makers risk either endorsing overly weak practices or demanding levels of protection that are misaligned with industry realities. Furthermore, we acknowledge that policy makers have a set of incentives that require them to be comparatively more risk averse—they are charged with safety, about which caution is warranted.

This makes the DP Deployment Registry, with its exemplars of existing deployments, valuable to policy makers as concrete evidence with which to benchmark new systems. By offering comparisons across implementations, it enables the public to ask why protections are absent when they have been adopted by similar systems. By describing the concrete fields that define a system's choices, the DP Deployment Registry allows policy makers to frame more informed questions, assess compliance with greater consistency, and refine guidance on the basis of real-world practice rather than anecdotal or fragmented sources. In this way, it provides both evidence for guidance and a practical tool for assessing compliance.

Moreover, anticipating this process will aid industry's own internal legal and policy reviewers, as they can more concretely anticipate the questions they may face and ground their own privacy bars in comparable deployments. This transparency strengthens accountability and creates a feedback loop between policy makers and practitioners: policy makers use the DP Deployment Registry to sharpen oversight and update guidance, while organizations gain clearer incentives to align with state-of-the-art practice and disclose their choices openly.

Taken together, the DP Deployment Registry supports policy makers across several dimensions: policy responsiveness, by ensuring that guidance keeps pace with evolving practice; benchmarking, by showing how new deployments compare to the state of the art; and evidence for guidance, by grounding oversight in the realities of deployment. By bringing these together, the DP Deployment Registry gives policy makers a practical lever to set clearer expectations, demand stronger protections, and push the entire field toward higher standards.

1.3. Value to Researchers

In addition to aiding deployments and fostering industry norms, we also envision that this information will be of use to researchers who wish to study the evolving privacy landscape of deployments. Such a resource might help researchers locate and analyze deployments of particular interest. At a higher-level, we also believe it will enable and encourage research studying ecosystem-level patterns. The DP Deployment Registry can also help identify research gaps, highlighting where real-world deployments struggle, and pointing toward open challenges in approaches and design choices. It will further support work that tracks the evolution of practice over time, allowing researchers to observe how parameter choices and techniques shift as the field matures.

Beyond analysis, the DP Deployment Registry can serve as a catalyst for collaboration and education. It may help researchers connect with organizations and practitioners deploying differential privacy to pursue joint studies or case-based research. It also has strong educational value, offering concrete, real-world examples that can be used in teaching about privacy. Broadly, because such a site will consolidate information that is already "public"

but distributed across numerous, long, technical sources, we imagine that this resource will aid and speed the study of practical privacy deployments.

2. Governance

NIST Public Working Group (PWG). NIST proposes to host the DP Deployment Registry through a Public Working Group (PWG), a flexible structure open to any member of the public. The PWG may establish subcommittees to perform work; for the DP Deployment Registry, we propose two: a Steering Subcommittee (process stewardship, outreach, sustainability) and an Editorial Subcommittee (schema stewardship and content adjudication). The PWG does not provide recommendations or consensus advice to the government and is organized to operate outside the scope of the Federal Advisory Committee Act (FACA) in both reality and appearance. NIST retains ultimate authority over the DP Deployment Registry, including the right to modify processes and content. The formation of the working group will be consistent with standard NIST policies of governing NIST public working group. Nonetheless, community effort through the PWG will be essential to ensure the quality and maximize the utility of the DP Deployment Registry.

2.1. Membership and Governance Structure

2.1.1. Bootstrapping Phase

In line with typical working procedures, NIST will appoint one or more NIST employees as co-chairs. Co-chairs would then post a public announcement calling for outside co-chairs to serve as an **Interim Steering Committee (ISC)** and any member of the public to join the working group. The Members of the ISC shall serve for no more than two years. To establish staggered terms, approximately half of the initial seats shall conclude after one year, determined by lot.

2.1.2. Ongoing Governance

After the bootstrap period, we expect the governance will transition to two committees under a public working group framework:

- **Steering Committee:** A body of **no fewer than three** and no more than nine experienced thought leaders and practitioners responsible for strategic guidance, positioning the DP Deployment Registry for community impact, and soliciting new contributions. Members shall serve **two-year terms**, renewable twice, with terms staggered so that approximately half the seats are contested each year. The Steering Committee sets broader expectations, ensures adherence to submission processes, and reviews Editorial Board recommendations. Decisions shall be consensus-seeking; failing consensus, a motion passes with a two-thirds majority. A quorum of **two-thirds of current members**, is required for any decision.

- **Editorial Board:** A body of **no fewer than three** and no more than nine subject-matter experts responsible for evaluating individual submissions to the DP Deployment Registry, maintaining the schema, and applying references such as NIST SP 800-226 [3] when interpreting privacy parameters. Members shall serve **two-year terms**, renewable twice, with staggering aligned to ensure continuity. Editorial Board members must disclose conflicts of interest and recuse themselves from evaluation where relevant. Decisions shall be consensus-seeking; failing consensus, a motion passes with a two-thirds majority *and* affirmative votes from at least four stakeholder groups. A quorum of **two-thirds of current members, with a minimum of six**, is required for any decision.

2.1.3. Stakeholder Representation

Both the Steering Committee and the Editorial Board would seek to include representation from across a wide group of stakeholders, including:

1. **Academics**: Experts in differential privacy, cryptography, data privacy, and data ethics.
2. **Practitioner Bodies**: Organizations such as NIST or recognized professional societies.
3. **Industry Leaders**: Representatives from sectors implementing differential privacy in real-world settings.
4. **Practicing Data Protection Lawyers**: Professionals providing legal and risk guidance.
5. **Regulators**: Representatives from data protection authorities or similar agencies.

2.2. Term Limits and Rotation

- Members of both the Steering Committee and Editorial Board shall serve **two-year terms**, renewable twice.
- No individual may serve more than three consecutive terms on the same body.
- After completing three consecutive terms, a member must rotate off for at least one full term (two years) before being eligible for re-election or reappointment.
- Terms shall be staggered so that approximately half of the seats on each body are contested every year.

2.2.1. Transparency

All membership rolls, term expirations, meeting agendas, minutes, votes, and rationales shall be published openly in the collaboration space.

2.3. Data Dissemination and Preservation Goals

NIST authority. The DP Deployment Registry is a NIST resource; NIST retains the right to modify, correct, or remove the DP Deployment Registry content and to adjust dissemination cadence and channels.

The DP Deployment Registry seeks to support broad and durable access to registry information for the research and practice communities. To this end, the DP Deployment Registry will regularly disseminate Public Data Files through a community data archive (described below) comprising all registry records, under an open license. The DP Deployment Registry aims for these Public Data Files to be provided in a manner consistent with the FAIR Data Principles – enabling computational systems to find, access, interoperate, and reuse DP Deployment Registry data with no human intervention [4].

As a complement, more frequently updated Data may be made available through other channels to support interoperation with other services, by approval of the Steering Committee. These releases may be made available as replica databases, snapshots of the early-release Public Data Files, feeds of updates, API calls, or other mechanisms. These complementary release channels are not subject to any of the terms of this section.

NIST may revise these goals and associated practices as needed to ensure consistency with NIST policy and mission.

Data Coverage

All records comprising the DP Deployment Registry database, in the core schema (as described above) that are used to provide DP Deployment Registry services.

Excluded Data

Except as later designated by the Steering Committee, registry website content and branding; the content of discussion lists; publications; and other supplementary documentation will not be included or as part of the regular public data release.

Documentation, Metadata and Bibliographic Information

Public Data Files releases will include metadata detailing the sources, review, and approval of the data—as consistent with the review processes described above. The release will also include administrative metadata used related to these records that are used to provide DP Deployment Registry services—such as modification dates, and any annotations on the data stored in the DP Deployment Registry service itself; any other metadata required to interoperate or reuse released data; any human readable documentation for the release that are required, in addition to the metadata, to meaningfully reuse the data.

Public Data Availability

NIST proposes to make biannual data releases to https://data.nist.gov, an archival data repository designed with FAIR principles [4] and best practice for Federal Data Strategy. The server supports persistent identifiers, data citation, automated discovery, and machine access. A NIST-owned GitHub repository will host the working copy of the data and be used to accept additions and amendments to the DP Deployment Registry. The GitHub repository will be used to release periodic stable versions of the database aiming for quarterly updates.

Audience

The primary intended audience comprises privacy researchers, policy-makers, and industry practitioners – who may use this data for research, policy, or commercial purposes. The data is also intended to be accessible to the public media community and the broader public.

Formats

Consistent with NIST data practices, all data will be made available in non-proprietary, community-recognized, machine-actionable formats.

Versioning

The DP Deployment Registry will publish updates to the Public Data File through a NIST GitHub repository. We propose to retain the entire change log throughout the life of the DP Deployment Registry.

Archiving, Preservation, Long-term Access

The DP Deployment Registry aims to ensure durable access to the Public Data File for research and policy. NIST will maintain data in line with the NIST Data Plan and will fall under the NIST public domain software license. We encourage the community to mirror the data, especially in credible permanent public access databases (e.g., Zenodo, Harvard Dataverse, Figshare).

Privacy and Intellectual Property

Consistent with NIST expectations for public working groups, submissions to the database must be non-proprietary. All information submitted to the registry must not include personally-identifiable information.

2.4. Policies and Licenses

2.4.1. Code of conduct

Communications with the DP Deployment Registry shall be subject to the following code of conduct, which may be modified in the future with the approval of the Steering Committee.

Introduction

This code of conduct applies to all spaces the DP Deployment Registry manages, including public and private mailing lists, issue trackers, wikis, blogs, Twitter, and any other communication channels our communities use.

We expect everyone who participates in the DP Deployment Registry community formally or informally, or claims any affiliation with the DP Deployment Registry, in any registry-related activities, and especially when representing the DP Deployment Registry in any role to honor this code of conduct.

This code is not exhaustive or complete. It distills our common understanding of a collaborative, shared environment and goals. We expect all members of the DP Deployment Registry community to follow it in spirit as much as in the letter, so that it can enrich all of us and the technical communities in which we participate.

Specific Guidelines

- Be open. We invite anyone to participate in our community. We prefer to use public methods of communication for project-related messages, unless discussing something sensitive.
- Be kind and behave professionally. Be inclusive of people from different backgrounds and of different ways of thinking.
- Be empathetic, welcoming, friendly, and patient. We work together to resolve conflicts, assume good intentions, and do our best to act empathetically.
- Be collaborative. Other people will use our work, and we, in turn, depend on the work of others.
- Be inquisitive. Those who receive a question should be responsive and helpful, within the context of our shared goal of improving the DP Deployment Registry.
- Be careful in the words that we choose. Whether we are participating as professionals or volunteers, we value professionalism in all interactions, and take responsibility for our own speech. Do not insult or put down other participants. Harassment and other exclusionary behavior are not acceptable.

continued on next page

continued from previous page

Reporting Guidelines

While all participants should adhere to this code of conduct, we recognize that sometimes people may have a bad day, or may be unaware of some of the code's guidelines. When that happens, you may reply to them and point out this code of conduct. Such messages may be in public or in private, whatever is most appropriate. However, regardless of whether the message is public or not, it should still adhere to the relevant parts of this code of conduct; in particular, it should not be abusive or disrespectful.

Assume good faith; it is more likely that participants are unaware of their bad behavior than that they intentionally try to degrade the quality of the discussion. Should there be difficulties in dealing with the situation, you may report your compliance issues in confidence to the NIST working group chair.

Acknowledgments

This statement draws on the Apache Foundation Code of Conduct for inspiration.

2.4.2. Privacy Policy

The DP Deployment Registry seeks to support broad and durable access to registry information for the research and practice communities. To that end, the DP Deployment Registry regularly releases a Public Data File, comprising all registry records, under the NIST public domain software license.

We remind readers that NIST does not own the rights to resources that are listed by links in a DP Deployment Registry record, for example, journal articles.

The NIST website privacy policy will apply to interactions with the DP Deployment Registry. Submissions to the registry use the NIST GitHub site. A name and email address will be required for any submission.

We request that you follow these community norms in using the Public Data File.

- Spread the Word: Please give attribution to the DP Deployment Registry as the source of the Public Data File.
- Be Fair and Lawful: Do not modify any data so as to make it false, incomplete, defamatory, or misleading.
- Keep us posted: We are very interested in hearing about ways in which people are making use of the Public Data File. Please contact us if you are willing to share your experiences.

3. Schema for the database

The core product of this endeavor is a database describing deployments of differential privacy. To make this information systematized and comparable, and to richly and fully describe the details that are necessary for meaningful understanding and comparison, we propose a schema of fields to be coded for each deployment. This schema is directly based on ongoing research by Priyanka Nanayakkara, Elena Ghazi, and Salil Vadhan [5], which in turn draws on [1, 3, 6–8]. We further emphasize that the schema will evolve as we learn what characteristics are important points of comparison.

Furthermore, the schema aims to cover all the dimensions that fully characterize a differentially private system, although this involves some degrees of depth that will not be equally meaningful for all deployments. The schema features structured fields that are directly comparable across entries, and the schema also intentionally carries many free text fields for compiling the details that meaningfully summarize additional implementation details from the published sources.

3.1. Transparency level of entries

The database has been initiated by collecting information from public sources, such as articles, white papers and industry blogs. The amount of available detail can differ substantially across deployed systems. As a broad-brush summary of the amount of available information about a system, we have divided the complete list into three tiers of information, from those that we would expect in even the sparsest high-level summary (Tier 1) to core parameters of the differential privacy guarantee (Tier 2) to nuanced implementation details (Tier 3). These tiers are not an indication of the quality of a deployment; rather, they convey the transparency level of the entry.

- Tier 1: Basic information (who is the curator, what is the intended use, a description of the system, a publication date, etc).
- Tier 2: Fundamental DP description (information in Tier 1 plus: DP variant, privacy unit, privacy parameters, deployment model, etc.).
- Tier 3: Nuanced deployment details (information in Tiers 1 and 2 plus: data domain, unprotected quantities, privacy accounting, implementation details, rationale, etc.).

The fields included in each tier are described below in Table 1.

488 **Table 1.** Version 1 of the Differential Privacy Deployment Schema. This is a living document that will evolve as more is learned.

Field Name	Description	Tier
Basic		
Name	Name of the privacy-protected data product	1
Data curator	The entity or entities publishing the data product	1
Description	Brief description of the product	1
Intended use	Intended use(s) of the data product	1
Data product type	Category of the data product, such as "summary statistics" or "machine learning model"	1
Data product region	Region the data product describes	1
Publication date	Date the data product was published. For publications with many releases, this is the publication date of the first release.	1
Data product sector	Sector(s) or domain(s) of the data product (e.g., technology, healthcare, education, government, energy)	1
Type of data curator	Category (or categories) that best describe the entity publishing the data product (e.g., government, industry, nonprofit)	1
DP variant		
Variant name	The variant of DP used, such as pure DP [9], approximate DP [10], zero-concentrated DP [11], and Rényi DP [12]. If a "custom" variant is used, then specify the full privacy relation, including the input metric, bound on input distance, output measure, and bound on output distance [6–8]	2
Data domain	Datasets "eligible for privacy protection"; in other words, "the actual, potential or counterfactual datasets that are to be protected" under DP [8]	3
Unprotected quantities	Quantities computed on the underlying data that are published without DP protections	3
Privacy loss		
Privacy unit	High-level description of the entity being protected, i.e., the entity whose data changes under adjacent datasets [8] (e.g., user, household, the set of all events associated with a user in a day ["user-day"]) [8]	2
Privacy unit description	Precise specification of what constitutes adjacent datasets in terms of the underlying data's schema	3
Privacy parameters	Privacy parameters (e.g., ε, δ, ρ). Which parameters are specified will vary according to the DP variant.	2
Privacy parameters details	Additional details or interpretations for parameters, if necessary	—
Deployment model		
Model name	Deployment model, such as central or local [9, 13]	2
Trust assumptions	Describe relevant actors (anyone who may access the data product or underlying data) and trust assumptions for each. Include rationale for these assumptions where possible.	3
Release type	Whether there is one release ("one release") or multiple over time ("many releases")	2
Release type details	For one release: Note whether there are future plans to release additional data products based on the underlying data. For many releases: Note the refreshment timeframe, how privacy loss is managed over time, and whether a fixed amount of privacy loss is allowed before the underlying data are no longer queried.	3
Data source	"Dynamic" or "Static." If the underlying data are dynamic, it means that new underlying data will be added over time. On the other hand, if the underlying data are static, new data will not be added over time.	2
Access type	"Interactive" or "Non-interactive." Under interactive deployments, people with permission (e.g., data analysts), can construct sequences of new queries on the underlying data under DP, including in response to previous DP queries. In non-interactive systems the queries are all defined before any DP release or privacy-loss budget is used. Examples of non-interactive deployments include synthetic datasets and aggregate statistics published with DP protections, but also algorithms where the desired summary statistics are known in advance.	2

Continued on next page

Table 1 (continued)

Field Name	Description	Tier
Access type details	If interactive, describe how the privacy loss budget is apportioned to and across people with interactive access, and whether non-collusion between said people is assumed.	3
Implementation		
Preprocessing and hyperparameter tuning	Description of preprocessing, exploratory data analysis on the underlying data, hyperparameter tuning, and how privacy loss from those stages was tracked, if at all	3
Mechanisms	Mechanisms used (Laplace, Gaussian, etc.), libraries, implementation details. Also includes: protections against timing and floating point attacks [14], whether steps were taken to prevent against leakage due to idiosyncrasies of the computing platform, and code/GitHub links if available	3
Composition	How privacy loss is accounted across multiple differentially private queries (e.g., sequential or parallel composition)	3
Post processing	Functions applied to the data product after it has been protected under DP	3
Justification	Justification and process by which any of the above choices surrounding implementation of DP were made, and well as any rationale around these decisions. Some questions that this section may answer include, but are not limited to: "What were the assumptions, modelling decisions, thresholds, and subjective decisions made in determining [implementation choices]? Why is the approach a thorough test of the stated assumptions? Was the process validated and verified? If so, how?" [1]	3
Additional information	Links to white papers, blog posts, or other resources describing the deployment. May also include links to the data product if publicly available, and any other miscellaneous notes that do not fit into any of the above sections.	1
Administrative		
Status	Approval status of this entry. Possible states are: Draft, Pending, Changes Required, Approved, Approved (Update Requested), Approved (Pending)	NA
Registry authors	The author(s) of this record. This refers to who filled out the entry, which might be different from who made the deployment.	NA
Tier	The transparency level of the description (either 1, 2, or 3). Higher values indicate more information is provided.	NA

4. Contribution process

The registry begins with publicly documented deployments, but its value will grow with community contributions. Organizations should be able to submit entries detailing their DP implementations and updates to existing ones. In addition, organizations should be able to directly cite and link from registry entry to registry entry—allowing for easy browsing from inspiration to additional deployments.

Proposed initial rules for publicly documented deployments added by the community:

1. Submissions must rely on publicly available documents, ideally primary sources.
2. Submissions should be based on public disclosures from organizations and should include clear citations and a link to the organizational website describing the case study. Published case studies must link back to their own registry entry as a form of verification.
3. Organizations can submit corrections or updates to their entries.

Proposed initial rules for deployments self-reported by organizations:

1. Organizations self-reporting their DP deployments act as a "promise" of their practices at a given moment in time.
2. An optional discussion board will enable the community to provide feedback and discuss registry entries, fostering transparency and accountability without formal vetting of technical claims.
3. Final submissions and corrections will be assessed by an editorial board to ensure their completeness, plausibility, and accuracy.
4. No deployments based solely on informal or second-hand information (e.g., hearsay, unverified third-party reporting) will be included. However, organizations may self-report deployments through an authorized attestation (see TB-3), even if they have not otherwise made a separate public announcement.

To streamline this process, we plan to manage submissions, reviews, and merges through GitHub Issues—allowing for transparency and accessibility. In particular, we will use Issue Templates that include a suite of questions based on the National Science Foundation's Research Coordination Network's template for PETs implementation use case guides. This template has mandatory and optional sections, however it supports maintaining a structured set of inbound information. In addition, we will create separate details for public case studies versus use cases where parameterization is not public. Beneficially, it also enables clear attribution of the work that volunteers contribute to the project.

4.1. Evidence & Trust Basis

The DP Deployment Registry is disclosure-oriented. The goal is to make more information public, and organize that which is public already. To that end, the DP Deployment Registry focuses on ensuring that every asserted field in an entry must be supported by a Trust Basis that is recorded and displayed. However, all of the trust bases are, essentially, attestations; it is not envisioned that the DP Deployment Registry will conduct code reviews or other forms of in-depth investigation.

The initial list of trust bases are listed below, though additional bases may be recommended by the Editorial committee or PWG co-chairs.

TB-1 Peer-reviewed publication (citation link)

TB-2 Formal publication by the deploying organization (e.g., white paper, technical report, product documentation)

TB-3 Authorized organizational attestation (submission by an identified representative using an official domain, with an explicit attestation statement)

TB-4 Other public source (e.g., conference presentation, blog post) — subject to heightened plausibility review

Minimum evidence thresholds. Hearsay, unverified third-party reporting without primary sources, and purely promotional materials are not accepted.

Editorial checks. Editorial review confirms semantic and syntactic completeness, internal consistency, and provenance clarity; editors do not assert correctness beyond minimal plausibility checks.

Verification. Specific fields or an entire entry may be labeled Verified when two Editorial reviewers have confirmed the relevant claim(s) directly against TB-1/TB-2/TB-3/TB-4 materials, or via a structured interview with an authorized representative.

Transparency. Each field's Trust Basis is recorded, and the public interface supports filtering by Trust Basis and by Transparency Level (see §3.1). Where information is missing or relies on TB-4, this is explicitly shown.

4.2. Orientation for record submitters

New data coders—by which we mean individuals who categorize the unstructured sources into the fields of the schema—will undergo a standardized training process to familiarize themselves with the schema, run by editorial board members and experienced data coders. This includes discussion of the goals of the project so as to orient coders to look for the correct details and provide the right depth of information. This may include example exercises from a previously recorded source. In the event a curator would like to add their information to the Registry but does not have the capacity to train their own personnel through orientation, they can collaborate with a trained coder from the project to assist in recording details following the practices of other records.

5. Editorial process for adjudicating contributions

NIST proposes to establish a public working group to assist in maintaining and growing the database. From the members of the working group, we propose to create two subcommittees: an editorial board that judges individual data submissions in detail, and a steering committee that sets broader policies and positions the strategy of the project to achieve impact. A draft charter summary for the PWG is provided in §2 (Governance).

Editorial Board: Some individuals would participate in an editorial board that would have authority to make decisions about the individual items and representations in the registry. For example, how to solicit adding new deployments to the registry, how to describe deployments whose details have been published, and what information and schema to use to generalize this knowledge.

Steering Committee: A small number of experienced thought leaders and practitioners who can contribute guidance for strategic decisions, and ideas for engaging the community of practitioners and fostering industry norms. The committee ensures that the registry

maintains its integrity, adheres to submission policies, and evolves to meet the community's needs.

Both of these committees will benefit from having members representing:

1. Academics: Experts in differential privacy, cryptography, data privacy and data ethics.
2. Practitioner Bodies: Organizations such as NIST that set privacy standards.
3. Industry Leaders: Representatives from sectors implementing differential privacy in real world settings.
4. Practicing Data Protection Lawyers: Representatives from the professional sphere who provide guidance and risk analysis to organizations as they make their data public.
5. Regulators: Representatives from working data protection authorities who are tasked with generating guidance for the use of privacy enhancing technologies in research and commerce.

5.1. Submission process

A new entry recording the attributes of a deployment will move through a sequence of status states as it passes through this process. When a record has been merged into the Git repository it is given the administrative status of "Draft". When the submitter believes the record is complete and ready for review of the Editorial Board, they can update the status to "Pending". If the Editorial Board approves the addition, its status is upgraded to "Approved", otherwise it can be sent back to any submitter with the status "Changes Required."

Only approved submissions are publicly shown on the NIST portal, however all GitHub issues are public and browsable from within the repository.

If the Editorial Board has reason to believe a deployment has changed or a part of an entry is in question, they can update the status of an approved entry to either "Approved (Update Requested)" in which case the entry remains on the portal, or "Changes Required" which reverts it to unpublished status. Both of these states have a Pending status that means the submission is revised and ready for review of the Editorial Board.

If revisions are made, direct links are provided on the NIST portal to the diffs in the Git repository.

This workflow process is visualized in Figure 1.

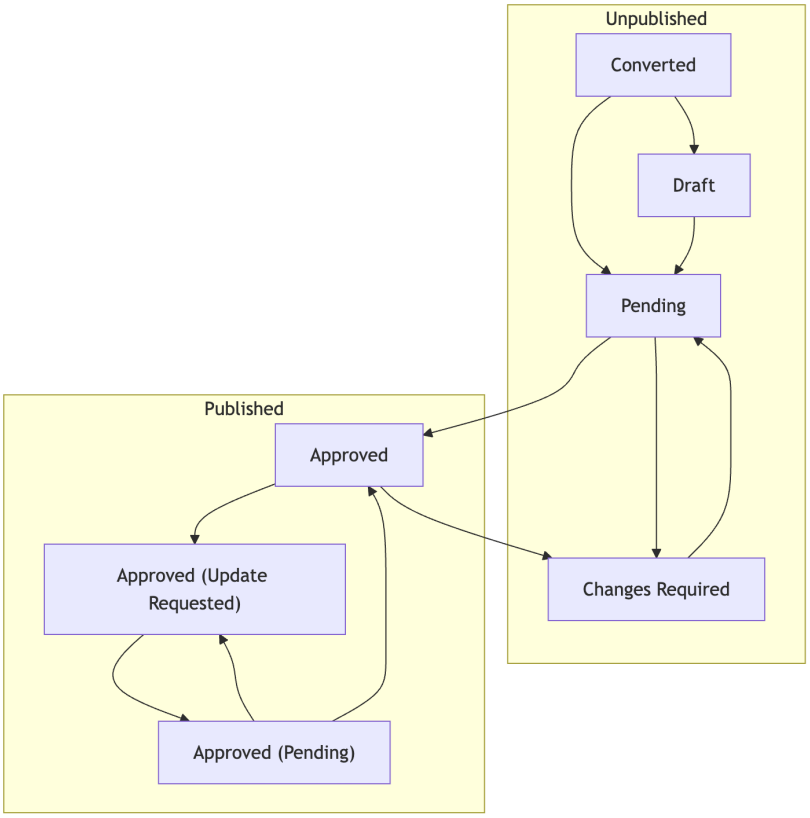

Fig. 1. Overview of stages a submission may pass through from initial coding to acceptance and any further updates.

6. Interface portal overview

We have constructed a prototype interface web portal to facilitate easy exploration and use of the data contained in GitHub, as well as sufficient materials to explain the information to potential audiences of differing levels of expertise. A screenshot of this interface is shown in Figure 2. The core element is a table with brief summaries of all deployment records. Any element is clickable, which brings up a card with the full details for that deployment by the proposed schema (as seen to the right). Simple visualizations (such as bar charts of the number of deployments by various schema field categories, and the distribution over time) are offered. The schema itself, can also be accessed from this portal.

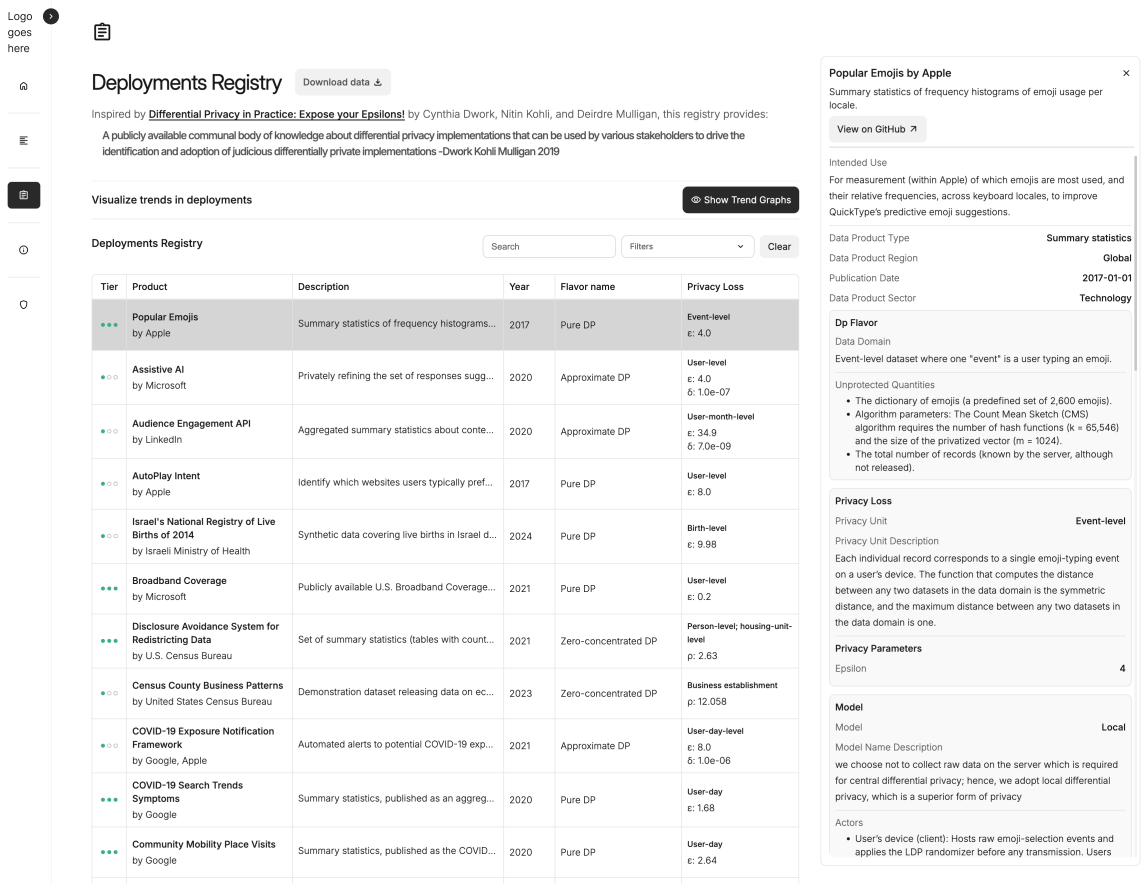

Fig. 2. Example of registry entry from prototype web interface.

Aspects of this design were based on ongoing research by Priyanka Nanayakkara, Elena Ghazi, and Salil Vadhan [5].

References

[1] Dwork C, Kohli N, Mulligan D (2019) Differential privacy in practice: Expose your epsilons! *Journal of Privacy and Confidentiality* 9(2). Available at https://escholarship.org/content/qt8vj5q8j7/qt8vj5q8j7_noSplash_3fbe06d219d4266437a90862727e5789.pdf.

[2] Desfontaines D (2021) A list of real-world uses of differential privacy, https://desfontain.es/blog/real-world-differential-privacy.html. Ted is writing things (personal blog).

[3] Near JP, Darais D, Lefkovitz N, Howarth GS (2025) Guidelines for evaluating differential privacy guarantees (National Institute of Standards and Technology, Gaithersburg, MD), 800-226. https://doi.org/10.6028/NIST.SP.800-226. Available at https://doi.org/10.6028/NIST.SP.800-226

[4] Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, Blomberg N, Boiten JW, da Silva Santos LB, Bourne PE, et al. (2016) The fair guiding principles for scientific data management and stewardship. *Scientific data* 3(1):1–9. https://doi.org/https://doi.org/10.1038/sdata.2016.18

[5] Nanayakkara P, Ghazi E, Vadhan S (2025) Unpublished manuscript. Available at url-unavailable.

[6] Cowan E, Shoemate M, Pereira M (2024) *Hands-On Differential Privacy* (O'Reilly Media, Inc.). Available at https://www.oreilly.com/library/view/hands-on-differential-privacy/9781492097730/.

[7] Gaboardi M, Hay M, Vadhan S (2020) A programming framework for OpenDP. Available at https://projects.iq.harvard.edu/files/opendp/files/opendp_programming_framework_11may2020_1_01.pdf.

[8] Bailie J (2025) *Topics in Privacy, Data Privacy and Differential Privacy*. Ph.D. thesis. Harvard University, . Available at https://search.proquest.com/openview/884fa8c618269402803855d40949fa72/1.

[9] Dwork C, McSherry F, Nissim K, Smith A (2006) Calibrating noise to sensitivity in private data analysis. *Theory of cryptography conference* (Springer), pp 265–284. Available at https://uvammm.github.io/docs/dwork.pdf.

[10] Dwork C, Kenthapadi K, McSherry F, Mironov I, Naor M (2006) Our data, ourselves: Privacy via distributed noise generation. *Annual international conference on the theory and applications of cryptographic techniques* (Springer), pp 486–503. Available at https://ptolemy.berkeley.edu/projects/truststc/pubs/101/ourDataOurselves.pdf.

[11] Bun M, Steinke T (2016) Concentrated differential privacy: Simplifications, extensions, and lower bounds. *Theory of cryptography conference* (Springer), pp 635–658. Available at https://arxiv.org/pdf/1605.02065.

[12] Mironov I (2017) Rényi differential privacy. *2017 IEEE 30th computer security foundations symposium (CSF)* (IEEE), pp 263–275. Available at https://ieeexplore.ieee.org/iel7/8048777/8049639/08049725.pdf.

[13] Kasiviswanathan SP, Lee HK, Nissim K, Raskhodnikova S, Smith A (2011) What can we learn privately? *SIAM Journal on Computing* 40(3):793–826. Available at https://epubs.siam.org/doi/pdf/10.1137/090756090.

[14] Jin J, McMurtry E, Rubinstein BI, Ohrimenko O (2022) Are we there yet? timing and floating-point attacks on differential privacy systems. *2022 IEEE Symposium on security and privacy (SP)* (IEEE), pp 473–488. Available at https://ieeexplore.ieee.org/iel7/9833550/9833558/09833672.pdf.