

# Differential Privacy at WMF

A case study

*Hal Tiedman – Privacy Engineer*

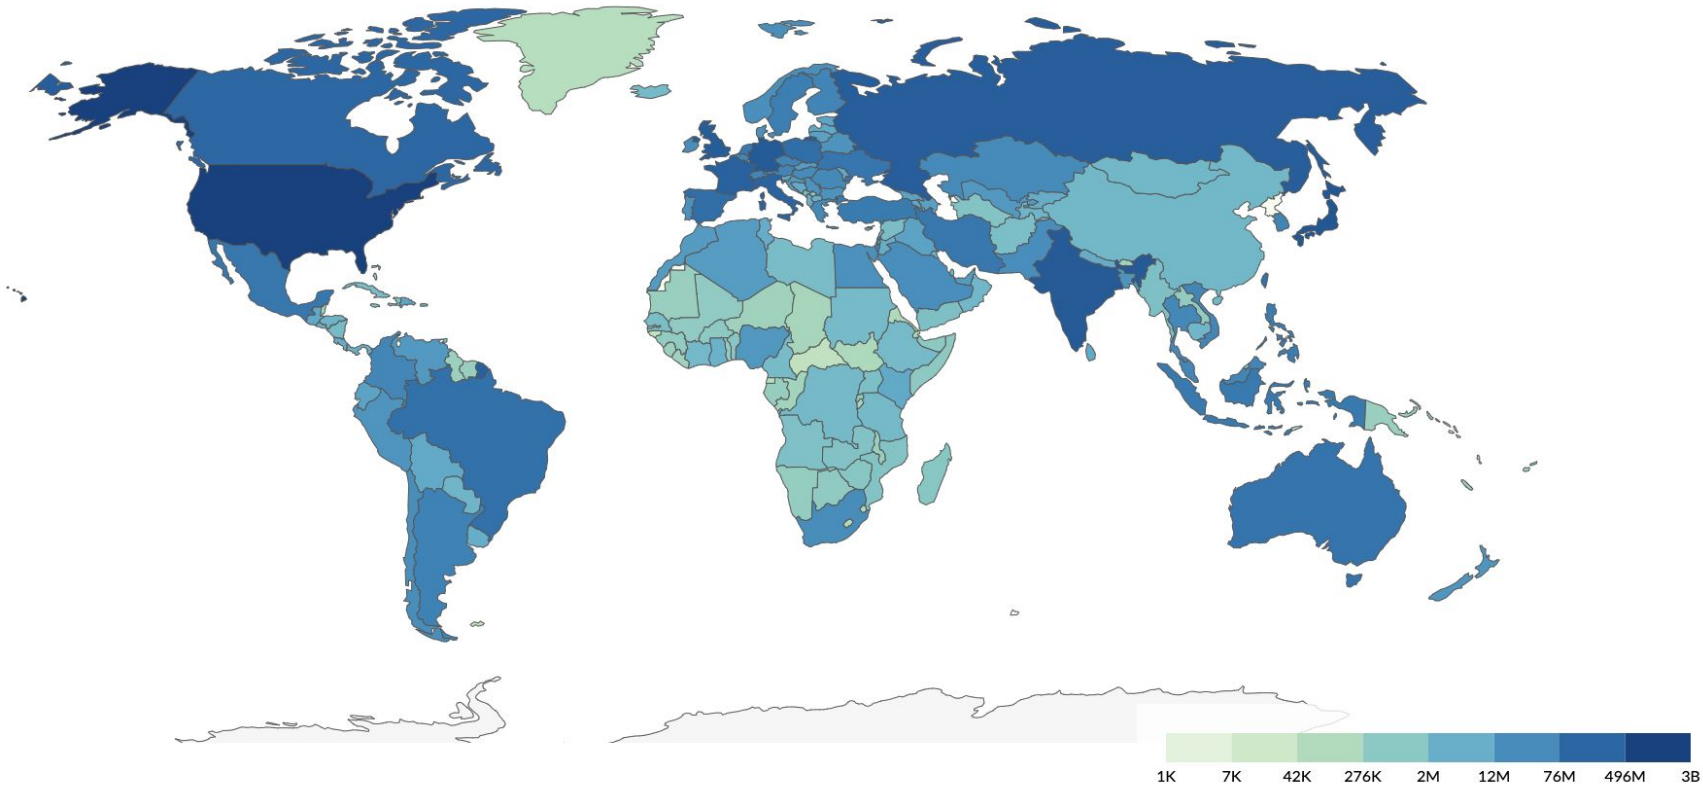


# What we'll cover

- Context
- What is differential privacy again?
- How are we using differential privacy at WMF?
- Why is differential privacy better than other options?
- Project roadmap

Context

# Page views by country



WMF wants to help editors and scientists understand user behavior by releasing data publicly (ex: Wikistats)

But what if we want to release more granular data?

Our foundational example: releasing  
pageview data by country and project

# Releasing pageview data by country and project

Could allow for useful  
disaggregations

Could also be dangerous  
for users


**Differential privacy (DP) is a good fit for this problem**

Wait... what is  
differential  
privacy again?



# What is DP?

A **process** takes a **database** in as input and returns some **data** as output



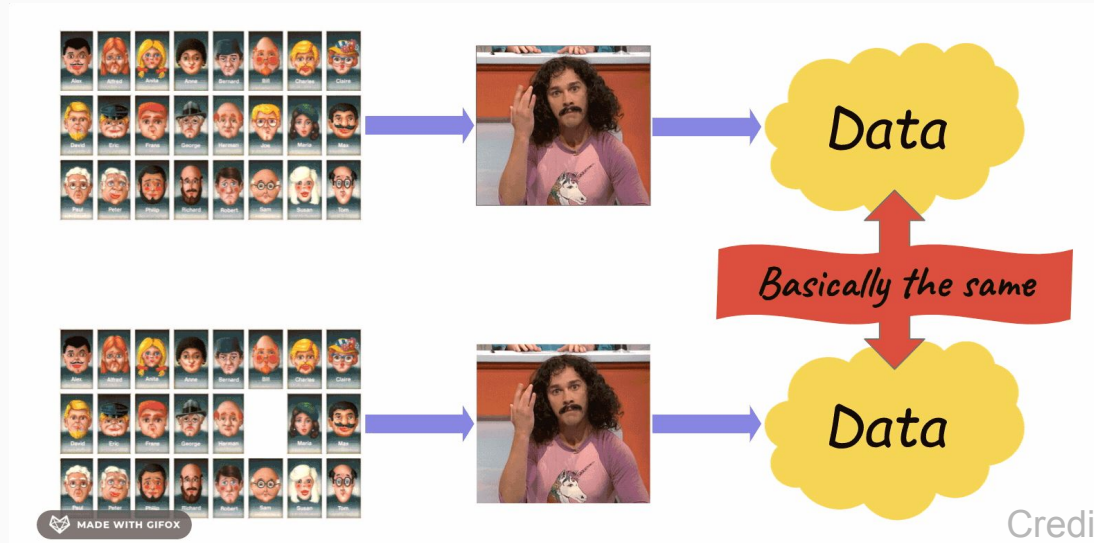
# What is DP?

- Add **random noise** (ignore for now how much, what type) to the process
- For now we'll call that **magic**



# What is DP?


- Remove someone from the database and re-run the process with **magic**
- Outputs should be **basically the same**



# What is DP?

**Basically the same:** Exact same outputs are possible with similar likelihood

Probability distribution **without** person in the database



Probability distribution **with** person in the database

# What is DP?

Differential privacy is a **promise** WMF can make to the readers and editors who contribute to our public releases:

From the perspective of someone looking at this data release, your contribution to this database will be hidden. High-level trends about the data will be visible, but no one will be able to infer your presence or absence in the data (even if you're an outlier).

# Differential privacy at WMF

# How are we using DP at WMF?



Credit: Wikimedia Commons

Why is differential  
privacy better than  
other options?




# Why is DP better than other options?

WMF's current approaches:



Filtering

+



Thresholding

+



Bucketing

Credit: Wikimedia Commons

# Why is DP better than other options?

WMF's current approaches:

- make it hard to define, quantify, or measure the privacy loss for a release
- force us to be overly conservative

# Why is DP better than other options?

Country	Number of page views for page X	Bucketed value	Included in the dataset?
<del>China</del>	1,371	N/A	No, filtered out
<del>Nigeria</del>	321	100-500	No below threshold
United States	10,262	>10,000	Yes

# Why is DP better than other options?

Initial results for releasing pageviews grouped by country and project:








# Why is DP better than other options?

- No assumptions about your attackers knowledge or capabilities
- Works no matter what an attacker already knows about your data

# Project roadmap

# Project roadmap

Project milestone	Timeframe	Done?
Background research into differential privacy and platforms	June – October 2021	
Outreach, contract, and onboard Tumult Labs	October 2021 – January 2022	
Install and test Tumult Labs software on analytics cluster	February – April 2022	
Design and implement algorithm for releasing pageviews partitioned by country and project	April – May 2022	
Productionize country-project-page-view algorithm	May – June 2022	
Run educational programming for WMF staff about DP	July – August 2022	<input type="checkbox"/>

# Interested in learning more about DP?

We're looking for:

- Technologists who want to experiment with this software
- More safe data releases
- Other potential applications for DP at WMF

Contact us using this Google Form: <https://forms.gle/f7sPCwjgUpQMoNSA6>



Questions?