Private Advertising Technology Community Group Minutes

Queue ←— Click Here to find QUEUE!

Scribes

- Day 1:
  - Scribe 1: bedfordsean
  - Scribe 2: npdoty
  - Scribe 3: stumakha
- Day 2:
  - Scribe 1: bedfordsean
  - Scribe 2: michaelkleber
  - Scribe 3: markkent

From the chairs - thanks!

Agenda

https://github.com/patcg/meetings/tree/main/2022/02/09-telecon
Webex Link: https://mit.webex.com/mit/j.php?MTID=mfd0fa61edf8ab8c3f4273dd4ea6631e8

Day 1

Administrivia

Introduction, Mode of Work, Resources, Participants, Agenda Check, and time for the various problems that come with getting on a webex and Google Doc

Aram and Sean Turner: Introductions, welcome, trying to get things on the table and get going. Wanted to go through a couple of things, don't want meetings every time to have a draft, want to use GitHub. Not sure if we need smaller subgroups but may do this if needed. People shouldn't feel like they're being left behind if they aren't here. We want editors of docs in GitHub to be active and move things along.
Sean Turner: Participants - would like observers, but recommend signing up on the W3C page so we can keep track of things.
Agenda bash - Does anyone want to add or change items on the agenda?
Nick Doty: There are qus in GitHub on charter for this community group. Do we plan to address this during the sessions?
Sean T: We should figure out how to address. Will play by ear and see if we have time for these sessions, or agenda for next time if not - also GitHub can be used for discussion.
Eric Rescorla: Didn't understand email about PCM discussion scheduled for the beginning of this. Concerned about having presentation/discussion coming from non member due to IPR implications.
Aram: Goal not to discuss active standards, but Apple have made representative available (John W) for any questions we have for today. They won't be doing a presentation/walkthrough of PCM.
Sean T: Other question we had was about recording thMOVe session. W3C rules are that we can record if no objections.

~Noted an objection therefore **NOT** recording the session
Aram: Note that nothing here is indicative of a particular product plan and discussion should not be taken as such unless it is explicitly stated
Aram: Any questions about WebEx, GDoc, or anything else before we go?.. OK lets go!

Charter for working group

- Working Group

Aram: Forming a charter around a Working Group. See a number of people already volunteering in this thread. This is the editorial group which is an outline to be brought back to larger group, then moved forward to the next step
Alex Cone: I volunteer, but do realise that several people with experience doing this is critical, so would like to hear if this should be people who are familiar rather than newbies?
Sean T: Ok if we have a mix of new people and those who have been around a while. We hope you will self organise around this. If you get stuck maybe defer those steps to those who have done it before.
Eric Rescorla: I liked what Sean just said. Happy to take a first cut but don't want to be "the boss". Happy to have someone else take the lead instead. I will do this unless someone else takes it. Would the chairs like this in a PR?
Aram: Have seen that CG for individual outputs sets up separate repos. If no objection, put in separate repos?
Brian May; Start it as a gdoc and move it once it has form?
Aram: Yes provided new editor is ok
Eric R: Yes ok.
Aram: Keep a link once you have it ready, via proposals repo, we can move it later in to its own repo space
Wendy: Wanted to offer context. Working groups are formally created after a review by the W3C Advisory Committee for their approval. Before that we send a work in progress notice ("Advance Notice") to W3C membership. Once we think it is at the right level of readiness to share, we want to give those with considerations or concerns to weigh in on inputs to the charter, so they don't have to file objections but can give input. W3C will take cues from the CG and chairs on when to make those moves.
Aram: Sounds good. We can move it to the next steps when ready
Martin: We have given Eric a "blank check" to write the charter, this is ok but we should circle back on this to ensure we have group discussion on what should be in the charter.Providing a little bit of input into that charter may be a good thing. After we discuss here we may have more to say.
Sean T & Eric R: Agree
Brad Lassey: Separate repo vs CG repo - suggest a separate repo since it can have issues tracked separately. Beyond initial cut, it should be in GitHub. Happy to help here too
Brian May: Think Martin's suggestion is a good one. What deadline do we want for the charter?
Aram: Early draft by next meeting?
Eric: I'm going to try to get this done in the next 1-2 weeks. Bother me if you haven't heard anything.

Privacy Principles Document

Aram: We form this document in our charter. Robin has already volunteered. We should probably loop back to it at end of day 2. This will exist to illustrate the concerns we believe we need to care about as this group

Robin: Happy to volunteer on this work. Ideally won't do it on my own. Better if other people join in probably. Not finalised yet but the W3C TAG (Technical Architecture Group) in conjunction with PING (Privacy Interest Group) are writing a doc of principles for the web and the idea is that this TAG document will apply to the entire web and apply to specific areas (e.g. Geolocation may have extensions on top of this principles doc for specific use case). One question I had is whether we want to focus this on privacy, which may be dealt with in the TAG doc, or whether we want to extend this document to include any principle we can stumble upon on how to do "web friendly advertising things".
Sean T: If you run into instances bigger than the original scope, note it in the doc. We can figure it out once we have something together.
Robin: Agree
Martin: Would like to thank Robin for taking this one. Believe it isn't as hard as he makes out - that work is in TAG/PING. Believe the purpose is to capture what we discussed on other topics as it relates to work done here. Don't expect this to have anything in it for a while, the role is to record important decisions that we make in the course of doing "real work"
Aram: Sounds good. It will be a living document most likely and we should consider it that way.
Nick: Curious about the potential other purposes we may have for this document? What are the advertising goals in a shared sense, how effective does an API need to be, what business cases need to be satisfied? E.g. if there's something like "Targeted advertising has to be **this** effective to produce an API", that would be useful to know. The web-adv group has use cases, not sure if we're planning to do this but it could be helpful.
Aram: Web-adv use case doc is a useful reference of comparison. When we come back to conversation end of tomorrow, the concepts and how they might work will give us a path towards something like this. Does that make sense?
Nick: Yes. Not sure if same document or another one.
Robin: Let's just write it and see where things land
James Rosewell: What policy decision would we take on such a document? What alignment will we have to laws or going beyond laws? Want to make sure the doc doesn't go outside a reasonable set of laws we would like to cover
Aram: We will note the laws that apply, but we don't want to be focussed on writing legislation or policy. Want to focus on technical ideas or limits for ideas going forward. Even that is vague because we are just at the start of this
Robin: This is quite simple, standards are by definition more restrictive than the law. If we define something that is illegal, people won't do it or they shouldn't. Just by nature those standards will be more restrictive.
Eric: In many cases the purpose of these standards is to protect privacy, for example we encrypt things even though US law prevents looking at data of individuals in most circumstances
Aram: I concur

- Privacy Principles

5m Break

Measurement Scope Discussion from @martinthomson

Aram: Martin asked to present on the scope he wanted to discuss for this group and for our conversation. Martin do you want to go ahead?

Martin: A few thoughts on strategy. PATCG can improve many things, including attribution, R&F, audience selection, retargeting, probably more. All of these are HARD. To be successful we need shared goals, understanding and trust, and focus, mutually beneficial for everyone.

Martin: If we try to do everything, we won't manage to do much of this. Suggest we start to concentrate clearly on a narrow set of goals. Suggest just one: Attribution. It comes up again and again in conversations

Martin: Our motivations for this; attribution is a major pain point. It benefits advertising of all types. It isn't so exposed to "hotter issues", and there are lots of potential solutions. Tracking is pretty terrible in terms of privacy.

Martin: We should do **nothing else for now** (in the order of the first few meetings/foreseeable future)**.**

Martin: Any questions/comments?

James Rosewell: Genuine question because I think it is important we understand the status quo. Can we define what the problem is with tracking and why it is terrible in terms of privacy to start with? This may help guide us to the solutions we can consider

Martin: This is something that is well documented at least from Mzoillas' perspective

Aram: Can put a pin in this question. Part of the intent of the rest of this meeting is for people to define how they see problems and how it has powered decisions

Sean T: You said "for now". CGs provide reports, and presumably you see we get to a point where the report is sufficiently stable?

Martin: Probably want to start on other topics before we reach that point, but want to at least get this on rails with people working on it, with a target, and we can then start preliminary discussions on other things without that discussion acting as a distraction. My concern is that we get so far into this and then start discussing FLEDGE, PARAKEET, etc and lose focus.

Sean T: Like the way of getting this ship moving in the right direction first.

Brian May: Think it is great to start with attribution, but it seems attribution cannot stand alone. We need other things like ads delivery, maintaining accounting on impressions shown, fraud concerns, so while attribution is a good starting point we will need to keep an eye out for things that need to be done in concert with attribution.

Martin: Good point, it's possible there's a v1 of attribution that will be very limited, but likely to evolve. I'd like to see us deliver something first, whatever it is and however limited it is

Brian May; Good to have something to break before we get started.

Andrew Pascoe: Brian covered my first point on the interconnected nature of these things. Also attribution benefits some channels more than others, for example things like search that tend to get attribution. A lot of the money comes from contextual too, so there could be an argument made for impression delivery mechanisms first since that covers more of the market

Martin: Haven't seen any proposals in that space

Andrew: FLEDGE and PARAKEET could be examples

Martin: Given size and complexity of those systems, I'd unfortunately suggest we start with something simpler

George London: Want to encourage us to be thoughtful and clear about the definition of attribution. Many of us have come from "direct response" advertising driving clicks/conversions/checkouts. Large portion is

focussed on driving attitudes, likelihood to take action, sentiment, etc Would caution us to consider this too as part of broader attribution definition.
Martin: Have been working on a definition. Not sure if it is a good one. If you want to refine we can work on it
James Rosewell: On queue to respond to answer earlier: Not many people from Europe on the call right now, so any decisions made may want to be tested over a period of time and when people from that geography can attend. Speaking more generally, if we take Mozilla's privacy policy, of which there are many URLs and docs that make sense at a high level, and going into detail, I would not personally agree. Suggest getting the areas of disagreement out at the beginning to sensibly debate. In UK the Information Commissioners Office (ICO) has published papers on e.g. pseudononymous identifiers and how these may work. Any approaches that remove e.g. all of these identifiers from the web and turns that into something else that only other participants in the web can use would be problematic if the statement was required under law and proposals ignored that elephant in the room. If we don't pay attention to that, we just create the same discussions under a different forum. My plea is that we can get to a place where we can discuss these subjects.
Aram: Thanks James, hopefully a lot of this can come out during discussion.
Brad: Practicality with this suggestion. Several other proposals that could fall under PATCG charter that are making good progress in other venues. Other folks want to move this all to one place, so practically it would mean this group focuses on attribution/measurement etc now and those other proposals would keep moving in other groups?
Martin: That is the plan
Aram: Technical note, it's up to those proposers to bring to this group when they are ready, and thus far they have not done so.
Brad: For Topics, it has been brought to this group as a proposal, should that be in WICG?
Aram: We will discuss tomorrow
Martin: Brief principles on attribution - learning whether actions on one site (e.g. showing an ad) produce results on another site. Have to start with a small concession; due to the nature of attribution, it necessarily involves the transfer of information about users, across sites, over time. This is a violation of privacy expectations.
Martin: Baseline assumptions, goal is default on mechanism with opt-out. The safeguards need to be strong for this reason. In our opinion, notice and consent is not an adequate safeguard.
Martin: Hard constraints, cannot and will not be able to eliminate bad actors. System has to provide privacy under those constraints.
Martin: Into opinion; not speaking to specific proposals, but likely aggregated systems could be acceptable. Those systems probably need to be distributed to prevent single entities having access to data about multiple users at the same time. We believe Multi Party Computation (MPC) is the most likely answer.
Martin: Want to start with something small - deliver something to demonstrate we can. Iterate from there to improve capabilities. This means that all of us need to be comfortable with getting something "less" than we wanted to begin with. This will still take much longer than we might like.
Aram: Thanks for the perspective, I see additional questions but will pause them as there are other topics to get to.

Private Ad Measurement and Attribution concepts

- https://github.com/patcg/meetings/issues/9

Aram: 3 sessions for today. First to chat with Wilander (Apple) to answer questions, not an official participant so not contributing ipr. John can take questions. Ben Savage (Meta/FB), and Charlie (Google) to follow.

PCM

JohnW: observers from Apple WebKit, working on joining CG. we have Private Click Measurement proposal in PrivacyCG / lots of overlap in membership. Proposal has been discussed for a couple of years. In beta 2019, shipped in 2021. More recently, adding cryptographic blind signature tokens as a way to fight some types of fraud. Working on alignment with Google on their Attribution API; request from TAG to align as much as possible.
JohnW: philosophy is client-side / staying on device. Reduce the amount of data combined between two sites, and sending that data in a way that is hard to tie back to a particular device (including IP address hiding, fuzzing timing).
Ekr: documentation on the cryptographic protections?
JohnW: "unlinkable tokens" on the click source side, when a measurable click happens, can request signing an unlinkable token, which will be included in the conversion, which may prove that it was a trustworthy click when it happened.

- Editor's Note: https://github.com/privacycg/private-click-measurement/issues/41 & https://webkit.org/blog/11940/pcm-click-fraud-prevention-and-attribution-sent-to-advertiser/

Ekr: so we're on the same page, what exactly is the use case? Payout, anti-fraud, training ML.
JohnW: about measuring clicks for incoming traffic. Not about viewthrough attribution or non-clicks. This limits the kinds of advertising that can be attributed. Only after-the-fact measurement. Not specific to advertising, can measure any kind of click with this. For example, could measure organic clicks, for a destination site to know where it's getting traffic from (even if not paid advertising). Navigational tracking also being discussed at privacycg, related to this work. Measuring click-based advertising.
Sean Bedford: entropy limited based on eTLD+1. For this group, would like to include both websites and apps. Where do we scope that entropy? Example.co.uk vs example.com might both be owned by the same company, but could potentially get twice the entropy out of PCM (or similar proposals).
JohnW: not designed to provide an advantage to a company that owns many domains. Etld+1 is a way to prevent subdomains from helping. Considering use cases where a domain has many merchants (e.g. Etsy): not limiting those cases as long as subdomains don't become a tracking vector for users all their own.
Sean: for companies with many brands, do we want them to share all that entropy across all their brands?
JohnW: for pcm, goal is to prevent cross-site tracking (not cross-organization or cross-company). This proposal sees example.com and example.co.uk as different sites.
Alex Brasil: seeing something like Fledge as a promising future direction. Would Apple/ios be interested in on-device attribution and auctions?
JohnW: no product announcements here. PCM is focused on attribution. But definitely following details of proposals in different groups, raising issues, etc. both bird name and otherwise.
Charlie Harrison: examples of partners who have successfully rolled out PCM and how it worked out, has that ever been documented as a fully-launched option? Data from the partner?
JohnW: we don't have that data; we do sometimes get requests/proposals about different use cases. We have blogged about some that we've heard.
@@link

Csharrison: if you could solicit feedback from people using it, it would be super compelling.
Robin: no promises, but I might be able to help with that.
Aram: us too.
Mariana: re: blind signatures, what is signed, and is it tied to a particular impression?
JohnW: source server can sign a token, material signed is not tied to an impression, … won't help against a rogue web engine. Source server can decide in the moment whether they think the click is real and decide to sign or not sign.
Ben Savage: design philosophy questions: not philosophically committed that only clicks be measured, just a practical way of starting. If it were possible to achieve privacy goals with impression attribution, you would be open to it.
JohnW: yes, absolutely.
Ben: design the API so that no one including Apple has the complete view of the world. Not opposed to multi-party computation or server-side aggregation?
JohnW: we thought we could design a way to address the use case without involving the server, and we prefer to keep it simple, make it easier for other browsers to implement and for developers to be able to test.
Ben: those are appealing features, but generally how do you feel about server-side aggregation?
JohnW: can only speak to the existing proposal for the moment.

IPA

https://docs.google.com/document/d/1KpdSKD8-Rn0bWPTu4UtK54ks0yv2j22pA5SrAD9av4s/edit#heading=h.f4x9f0nqv28x

Ben: high-level overview and design principles. Interoperable Private Attribution. Use case is designed to provide just basic aggregate measurement. This is aggregate, not event-level (in contrast to some others). Attribution is done within multi-party computation. Protecting against adversarial consumers: add differentially private noise managing a privacy budget, (user, site, time window).

Ben: basic aggregate measurement is something we would like to support. Believe there is consensus to add new APIs to the web platform to provide that. And tech exists to do that in a very private way. Given that at w3c we can only make progress in areas where we can find consensus.

Ben: believe that there isn't consensus to add new APIs for remarketing or interest-based advertising inferred from browsing behavior; less consensus or might get stuck working on those more advanced use cases.

Ben: build trust by shipping something together first, demonstrate that we can work together to reach consensus across organizations, and ship something. Will be easier to make progress on those other topics once we have demonstrated success in this area.

Ben: Chrome proposal for event-level reports, which include a 64-bit identifier associated with a particular click. Did this specific click lead to a conversion or not? API is designed to prevent record-linkage, may include adding noise (for differential privacy or plausible deniability). Reactions from Safari and Mozilla suggest unlikely to find consensus ("outside the consensus window") or that it's very contentious, even if that design effectively prevents record linkage and has some level of plausible deniability. Because event level reports tells you some information about what a specific user did on another site: “there’s a 75% chance that Ben bought something on shop.example”. This still tells you something about that specific person even if it’s noisy, and could be sensitive domains (like donations to a political campaign). Not a wide consensus that by default individual records should be revealed in that way.

Ben: aggregate measurement APIs seeing a lot more support/proposals from different actors.

Ben: Big new thing that we are proposing in IPA is doing attribution within the MPC. This is new and different: With PCM, anonymous attribution report sent from the device to the server, including metadata (site of etld+1, 256 values of source id, 16 values of trigger_data, rough time window, and rough IP range). Attributed on site could be one of an unlimited number of registerable domains, which an attacker could register for different users. Attacker could combine the registerable domain and the source_id to uniquely identify the user. Seems to me a fatal flaw of using low-entropy identifiers. And source_id 256 limit is already going to be very limiting to sites that have many vendors, for example.

Ben: re Chrome Aggregated Reporting API, attributes events on device and then send (after a delay) a report with some metadata. Again, metadata and timing could be used to reidentify the user.

Ben: goal for IPA is that receiving a report reveals no new information beyond that you called the API. you already knew that an ad was shown or that a conversion happened when you requested the report. By saving up a whole batch of reports, the MPC can calculate a count or sum. Believe it’s a better privacy property, and also higher utility because it won’t require as many random delays. A significant pain point for Facebook to have that delay in the system. This also doesn’t rely on IP blindness in order to provide user privacy. And doesn’t require limiting entropy of the source/trigger event. By moving attribution to MPC rather than on device, can improve privacy and utility. And can work for cross-device or cross-browser/app/UA. Possible to do cross-device attribution using browser sync mechanisms, but wouldn’t work between different browsers.

Ben: by having standardized source and trigger events, IPA can be interoperable between browsers/devices. Especially important for smart TVs, where I don’t expect anyone will click on an ad. Or in-app web views. Dramatically undercount the effectiveness of advertising if not attributing cross-device.

Ben: preference for differential privacy. Grain could be limited more severely compared to Chrome aggregate api. [wrapping up quickly so scribe may have missed a lot of that.]

James R: thanks! Regarding consensus about new APIs. Not clear what the problem is with existing APIs. Discussing some of the very largest companies, but not the majority of participants in this group, and those larger companies are engaging in more groups. At w3c, believe we should use existing functionality / lego bricks. Proposals from different browsers or gatekeepers tend to play to their functionality/advantage. Not clear why we need to make any change from existing APIs.

Ben: Facebook doesn’t operate a major web browser, but looking at large browser vendors to find the possibility of shipping an API across major browsers. Don’t want to waste time on proposals that won’t be shipped by major web browsers.

James: confused why notice and consent wouldn't be adequate. For default on, not sure who controls the defaults. Why use a multi-party compute solution?

Martin: We do not have the time to fully answer that question.

Aram: Agreed, please open an issue in the proposal space or on the issue thread.

Charlie: re: on-server mpc attribution, consider the cons:

1. Complexity of calculating attribution on the server, harder to operate on ciphertext/secret shares. Increases complexity quite a bit.
2. Scaling challenges: orders of magnitude more intensive
3. Data minimization: better if the servers have less information rather than more. Splitting of data via MPC is one protection, but would always be better to just not send the sensitive data off the device.
4. Data deletion: delay also has an advantage for the user that they can choose to delete something before it gets sent out. User might want to clear their browsing history after doing something especially sensitive, could delete a pending report so it isn't sent to an aggregate system.

Ben: We can take this up elsewhere as well. Yes, this is more complex. All pre-attr - higher volume. Yes, higher volume and cost, we attempted to design where cost falls to the caller of the API .. we need to evaluate economically feasible? With regard to data minimization: agree that we want to minimize data, but counter assertion- on device reports sends a different type of information. If servers have sensitive info- we failed. Delay potentially giving a chance to delete? Not convinced. If we need this capability- people should be able to turn off, no delay needed. We designed not to reveal sensitive info.

Moshe: How are you going to do cross device linking? Where a user sees the ad on one device and makes the purchase on another browser, how would you differentiate between xyz and .

Ben: I am going to direct to G doc. We use match key. If logged in on dev A. If user logged in the same service. We can set the same match key there is no get match key, it is set only. When event is generated, we just specify what match key you would like to use. No way to limit the match key's readability across reports, anyone could read it. If you invoke match key, it would enable anyone to use match key.

Moshe: authentication required?

Ben: A login service would set that match key and then any other website on the internet can leverage that match key once it is set.

Moshe: thank you

Brian May: it would be very helpful to have basic data flow included. Next is possibility of doing on device before doing on server. THere is a question who pays for, owns, is responsible for the server, differences in EU, US, different jurisdictions for data transfer. It can be difficult conversations

Ben: no time to respond now, see: https://docs.google.com/presentation/d/1NpQz0Wm73eEKw24V7B0yCjq4Tw2qPgeezhMfS0-P-TY/edit?usp=sharing

Bosko: tech proposal . How you think should be permissions set ? If there a relationship needed with match key providers

Ben: My worry is that if a match-key provider has technical ability to limit match-keys they set, that's too much power to give them. I want API not to give anyone a particular advantage, everyone has equal access to attribution and measurement. If you call setMatchKey, and you're example.com, anyone can use example.com as a MatchKey provider, to avoid structural advantages accruing to any match key providers.

Bosko: what would be incentive for match key providers?

Ben: benefits of more accurate measurement. Having competition adv or accuracy. I imaging accuracy..

Shivan: Martin mentioned opt out. If client who opted out should be indistinguishable from one who has. How not to penalize?

Ben: It will not be likely challenging to design

Martin: Consistent results for the same user on the same site. The same machinery is going to go pulling and it will look like it is working but no matches produced on the other side.

Shivan: good, let us make sure remains the case

EKR: Two questions. Can you talk about the status of the sidecar use case (Editor's note: meaning additional data besides the match key)?

Ben: in addition to matchkeys, is there anything else in this events? 8 bit integer. Sum them up to get conversion. Honest but curious model. We are trying to get to malicious threat model. If anyone with crypto experience can help , we are interested

EKR: Can this be used for click measurement, training ML models, etc. It is hard to understand .. We are trying to get off curve
What are the different use cases and different sliders?

Ben: we are taking an iterative approach, starting with simple use case, counting, not ML learning model. We are targeting- can you count total number of conversions attributed to a campaign. What is return on ads, lift measurement, this is a use case. Is it causal? Test and control groups. If attribution happens in MPC, we can send .. feature vector and a . I am not opposed to this, just need to get going with a simpler use case

Attribution reporting

Charlie: John Delaney is joining me.
Framefork- John is going to go over design.

John: attribution reporting API.
What is our definition? Does an event on one site cause an event on another site?
This is on device attribution model. Browser is responsible for creating events and sending reports to the parties. Last touch, priority based, other considerations in the logic on browser. All gets to AdTech . We took a diff approach on what gets sent. ROI, conv counts. How many conv. Were in geo - in aggregate report. Browser sends encrypted reports. You can later sum them together. AdTech ..
Event level report- detailed fine grain- for ML learning where you need precise data. When browser sends data, it is high entropy. These can be simple text, no encryption. How many report.- some of the controls.
For aggr reports- histogram. Series of buckets where value is some integer in the range. Key can contain arbiter. Data. You can cross geo vs. other features in flexible way. This is a bit of overlap, but subtle differences.
We done our best to make report work in concert with each other. 10000 foot overview. Recommend reading docs.
Charlie: add in terms of privacy, we are adding noise and also randomized response to protect privacy. I wan to go over . There is elephant in ther room- utility vs privacy camps. On privacy- delayed data. You need to wait until you have a batch. We are talking about reduced entropy. We need to put more report in less buckets. There is a spectrum across these camps.
Fine grained- we can learn a lot.
It will be difficult to get consensus. We need to try and experiment still.
Design consideration- where are we putting complexity ? Client or server? All the responsibility on client on one end vs .. on server in another camp
Foundational decision we need to make is where is the complexity in the system, client, server, hybrid somewhere else?
[slide: on-device or off-device]
I think that this decision inform other decision. This on device mechanisms provide easy compatibility. Cross -device.- big one.
Server side attribution provides a bunch of huge super powers as has been mentioned.
[slide: server-reliance and privacy efficiency]
When I say privacy efficiency- more . If we are less reliant on servers, even if we are using a fixed definition of privacy, we are able to achieve it but get less utility. You can ramp up privacy, but the data gets difficult to interpret.

On the other hand, if we rely on servers, we are able to get protections work. IPA, MaskedLark- some examples of this approach.

Mixnets- interesting compromise to create a simpler server architecture. You do not need complexity, but get component that you shuffle around.
THis one is about agility. We talked about MPC, but MPC designs will need to be custom fit to specific tasks. Let us say we need a histogram, we will likely need new crypto research. This is trade-off of MPC.

Running on trusted server- trust issues.
Execution environment- need to check specific instance is running the code you intent it to run. It would be useful to run a survey on: [???]

[slide: delegation & privacy budgeting: 2 axes delegation hard-easy; privacy budget independent or shared]
How easy is it to delegate?
I a lot of times you see that publishers will work with many 3rd parties on measurements. Need to make sure some of them are not lying. If you need to run queries, etc to verify, it is complicated. Only competed played can run something like that
If you are Nike, then there is a question, should they share privacy budget? It should be based on the site?
There is an issue of fairness risk. What if one party uses common budget up?
We want to have independent budget? Then we have privacy risk. They can collude and average results to de-noise results
No matter where you move on diagram, there are issues. There is a range of dragons here.. Will add slides later.
Questions?
Aram: Slides, please send them to github for us to include in minutes
Ekr: comment- I do not feel that using a Trusted Execution Environments are feasible at all. There is a lot of work on attacking, and it’s trivial with physical access. Side-channel attacks are not in the SGX threat model.

Charlie: Great point . The point of the flight is that these things are not foolproof. I would say it is a marginal improvement but not worseless

Ekr: we need low value to attacker
Ben S: I‘m uninterested in trusted execution envs for another reason. it seems that there is a limited supply of trusted hardware in the world. This approach will run across this problem

Nick: There were comments on potentially there being a single trade-off of privacy vs utility. It is worth having more discussion on the trade off. Privacy is more complicated with controls, etc, beyond simple trade off. If we have a way for users to understand how users got an ad. This is just something to consider- there are different types of privacy. Control/ vs participation.

Charlie: You are right Nick. I was gong to agree. Yes, there are many dimensions to privacy, not all are in conflict with utility. We should go in with eyes open that this is the case that either critical use cased not supported or privacy expectation are not supported.
User does not to delete a report in IPA- maybe too strong assumption, as an example

Kleber: I would like to make clear that the question how to establish trust on large component is what we are talking about here. One is MPC .. there are other approaches. To achieve trust that the system is doing what we expect it to do. The identity of who is running this is also a component here. GARUDA is an example of API where trust can be established across parties. Do not discount other option than MPC

Charlie: Yes, we should consider the whole spectrum. I also think we can compose this options together, not necessarily chose between two options

Raimundo: while I love MPC, I see advantages to alternatives. Can we assign on device APIs so that they are flexible and work with multiple solutions

Wendell: commercial perspective. Will do tomorrow.

Kleber: ok, will cover that tomorrow. Thank you all
Sean: Thank you for not drawing lines. Very optimistic on the next steps.
Aram: Thank you all in uncomfortable time zones.

[Editor's note: we will likely want to turn off the Webex chat to keep discussion in IRC for our next meeting.]

Day 2

Opening Review / Q&A

We are using irc.w3.org channel #patcg today for sidebar chat.
https://irc.w3.org/?channels=patcg
Note: IRC content will be preserved

Queue is down the bottom of this doc (above the two participants lists) if you need to add yourself to it

Aram: Welcome. Yesterday we overviewed charter, there are a few issues that need to be resolved. Sean T worked on this with Nick already. We talked about process (see above). Martin gave us an opinion of goals from the Mozilla perspective, we had a Q&A with John Wilander on Apple PCM. Ben Savage gave us an overview of the IPA proposal authored jointly by Mozilla and Meta. Slides all in the repository under the meetings folder. Any outstanding questions about yesterday or about our approach for today?
Alex: Some good discussion unfolding on GitHub under proposals. Would like to ask if possible that if you open an issue, please scope it and don't change the name of the issue. Because if you change the topic the comments no longer make sense. Especially aggravating to the opening commenter since it looks like they are a "crazy person"
Aram: Agree, if changing the title substantially changes what the context is, it is easy to open new issues, so close the old issue and open a new one. Make sure to open the issue clearly, state a clear problem, and be clear and specific in the request or the suggestion so that we can use issues effectively.
David Dabbs: For future meetings and for others, would suggest putting IRC and similar link in logistics section of the agenda doc (in GitHub)
Aram: Also calling out comment from IRC - don't have to close old issues if they are unclear. We only want issues being closed if they are resolved.
Alex: Want to make clear, I don't want to close things unresolved, just asking for consistency in discussion titles.

Aram: No other people on the queue, so want to give Wendell an opportunity to ask the question that tailed off yesterday.
Wendell: Deferred yesterday because it wasn't the right time. Making a plea for commercial sense of the proposals. The IPA proposal and the other Google and MS proposals have been campaigned heavily, this is a huge investment of time. For those of us who will be users, we are trying to figure out how to get involved. This manifests in terms of cost (financial, eng, maintenance and so forth). Want to accent that sense in this group for understanding. PRIO is to our understanding in commercial practice but unclear how expensive it is for adtech scale.
Aram: Any questions or anyone want to respond to Wendell?
(No)

Private Ad Measurement and Attribution concepts

- https://github.com/patcg/meetings/issues/9

@joelpf Masked LARK / 10m Q&A

Joel: I'm Joel Pfeiffer, work on MaskedLARK with collaborators at MS. Charles has presented before. Inspired by Aggregate conversions MPC flow from Google. MaskedLARK is what we believe to be an extension of that. Ideally we'd like to compute a gradient for an ML model.
Joel: Related work is the Google proposal for conversion reporting using sMPC for trusted mediator. MPC performs the aggregation and enforces DP in the process. Upon leaving browser, no individual entity has a complete picture of an individual record.
Joel: Pros; segregated helpers is more palatable than single trusted mediator. Cons; handles aggregation for some reporting needs but not modelling.
Joel: Goal of this is to expand the aggregation service towards an abstract DP enabled Map-reduce network. Browsers implement map, helpers implement reduce, ad servers are the consumers and data storage.
Joel: Second goal is to expand from single aggregation to a platform that enables various DP functions (aggregation, model training)
Joel: Will talk specifically about model training today. (Papers linked in slides, prototype also available).
Joel: Browsers represent the user interest, for view/click of ads and subsequent conversion. They are the only party with full info on users. They have to attribute conversion to clicks, insert fake records+masks, they have to insert localDP, they have to share secret values for training, send reports to ad server.
Joel: Helpers receive the encrypted reports from ad servers. They perform the core function - aggregation or model training. Also enforce community privacy constraints such as k-anon, global DP, privacy budgets
Joel: Ad server queries helpers to get aggr data when needed.
Joel: Example in mind was gradient computation; where features are observed but labels are secret. Core issues is that "differentiable models" optimise via Stochastic Gradient Descent. This is computed per sample so need features and labels together. This would reveal information to the helper. Can use MPC with masking to do this anyway.
Joel: Browsers generate lots of possible labels/features, but will add fake ones along to the true ones. Both helpers independently compute gradients for true+fake samples. Browsers send masks to the helpers to include in summation. That mask is constructed in such a way that the fake examples are removed from final summation

Joel: (Showing slide 13 walk through example). Note that because summation is bi-linear, can sum in any direction and it will still work.
Joel: More generally, can cast feature/label/model space, so other functions can potentially apply.
Joel: Taking simple example, can prove we can take this to many labels. This is detailed in the paper
Joel: Fake and real can be sent from the browser, but order needs to be randomised. For real values, need to be quantized, for non integer, they can be randomly rounded.
Joel: Threat models - high cardinality label space, ad network could just send an ID and collude. This is why we randomly quantize labels. Browser poisoning - invalid mask for example, mitigate with overhead for validity (SNIPs). Requirement of default label when expired - randomised ending times needed too. Lastly one helper could collude with ads service. This could expose sensitive features. Could add noise or do a cost sharing scheme.
Joel: Considering efficiency/utility vs privacy. At efficiency end, event conversion API, at other end bucketed secret shares. Masked LARK toward the privacy spectrum but not as far as fully bucketed. Another option could be masked event conversion API - more efficient but less private.Could also go the other way with secret share Masked LARK toward the privacy end. Much more overhead but it could be done.
Joel: Considered some other things with local DP test and gradient clipping for noise etc. Wrote a PyTorch interface, and have some experimental results that expect to publish online. At this will stop, (summary of above).
Eric Rescorla (EKR): 2 Questions: In slides you show that client submitted its own results and fake results and the mask washed them out. It implied you would spit out computation one client at a time? Should they be batched?
Joel: Yes it is expected to add DP noise.
Eric: LocalDP or GlobalDP?
Joel: Both
Eric: Seems like the DP to add seems like a lot. Will read paper. Seems like you need an enormous amount of data to train the model.
Joel: If there is anything wrong or any issue please let us know!
Eric: PRIO has description of how to do OLS with PRIO, can you compare that to this?
Joel: No conflict there as far as I know although not familiar with PRIO
Charlie: Can you go into more detail on how we plan to protect feature vectors from collusion by adding localDP?
Joel: LocalDP adds noise to a dense feature vector to prevent recovering it (directly fuzzing it)
Joel: Two things in there - local DP but if that isn’t satisfying we could also do secret share too. It is possible but also more expensive. We believe that sometimes you may want to make that tradeoff and sometimes you don’t. E.g. if everything is already 1P it doesn’t matter, but in other cases, maybe FLEDGE, you need other information.
Charlie: Wondering if you considered revealing the feature vector to the servers is less sensitive? E.g. if the feature vector will be unconditionally sent regardless of user action, would mean you could consider it fully public knowledge. Even if you didn’t do all of this you could add noise only on deviations from that, e.g. assume some fraction of unattributed or attributed conversions overall and will send some reports unconditionally in some cases. Seems like other approaches of local privacy?
Joel: Yes, would say you’re right but wouldn’t say we have an exact solution or when to do it. A trick of this is how to accurately simulate things. For example in a given timerange need to handle this. Open to ideas from ƒPothers as we are a bit stuck. I work with search more than display, which is mostly 1P, where we felt that

mask with event conversions API is enough. Could go crazy and do much more on top of that, giant computes on top of the feature vector. Liked the idea because maybe “plausible deniability” isn’t enough and then you just have lots of labels in lots of buckets that you don’t know what they are, but this doesn’t really matter because you knew it already.
Charlie: Yes you could argue that if you did that and put DP on top of it, then it would be better privacy. If you set the system up so that the collusion leaks nothing, then it could be stronger.
Joel: Google has a big search engine too right? We figured that this would make sense. Can take more of it offline if needed
Martin: Funny when Charlie said he understood. I didn’t. Can take that offline. You asked previously about PRIO. One of the core realisations with PRIO is that addition is pretty easy. Making sure the numbers aren’t bad is the hard part. I know we have talked about this in the past, but what are your plans to deal with malicious clients?
Joel: Most of the information lies with the masks. I’m an ML guy so not up to speed with crypto. Seems simple to verify the masks are valid in the space and lie in 0-1. Seems like it isn’t a hard problem but am not an expert. Malicious user can’t corrupt too much since the data is provided largely by trusted parties.
Martin: SNIPs may be interesting for you to explore
Wendell: Area of cost - do you have a sense of how costly/inefficient this scheme is for standard workloads? You work with search, but the sense is that anything done in these schemes (secret shares, masking ,etc) is going to have some factor of inefficiency. Would be good to understand what that is - would liek to understand a “cost per record” or “factor per record”. We have graphics from yesterday of positioning this on spectrum. One could imagine assigning some estimate of factor of load required to achieve results. Do you have this?
Joel: Remains to be done. Our perspective that this can be done but may be that our models are too big.
Wendell: When people say expensive they often mean 2^2^N - is this prohibitively expensive or more like $10 per record?
Joel; Intuition is this is not prohibitively expensive no.
Brian May: WOndering similar to what Wendell asked - but spectrum of adoptability. In terms of understanding this it is a sophisticated system that most won’t understand. Also lots of infra+eng involved. How can we make this commodity consumable?
Joel: Don’t know the exact answer to that question. New hire implemented general idea out of the box. We have the code out and people can try it, there’s publicly available data sets.
Alex: Consideration - a lot of the protocols implemented that make the internet run cost a lot of money, but it’s about motivations. If there’s enough value, then for better or worse the economic systems will figure out how to pay for them. Think back to other standards and protocols that themselves have become quite expensive but are creating a lot of value for the internet.

[scribenick: kleber]

@marianapr on Prio/DPFs - https://datatracker.ietf.org/doc/draft-gpew-priv-ppm/ / 10m Q&A

Mariana: Aggregate API with PPM presentation
…: This is sort of in the wrong order vs the previous presentation, since we already did ML and now we're just doing intro MPC aggregation

…: Was very happy to start with MPC as commonly-understood term yesterday. I'm in the I-love-MPC club! But just to be sure we all know what it is:
…:Encryption and Signatures protect data in transit and at rest. MPC is the next generation, protecting data while computing on it. MPC is about enabling computation of a very specific function on a particular dataset without revealing more than the output of the function. Could include cases where parts of the dataset are owned by different parties.
…: I look forward to this being part of the infrastructure of the internet.
…: For the aggregate API that Chrome proposed: Differentially private histogram. Each client contributes a value to some bucket of a differentially-private histogram that we're trying to compute.
…: Based on the local information on the device (attribution done locally), bucket computed on the device. Even the choice of bucket is private, since it might reveal on-device private signals. (If we don't need the bucket to be private, things are different of course.)
…: Prio Secret Sharing (Corrigan-Gibbs)
…: User's device takes input, splits into two cryptographic shares (neither one alone reveals any information), handed to two servers that can do aggregation
…: Includes cryptographic range proofs, so that the device can send a proof that the input it's providing is in some pre-specified range, so a malicious client can't mess up the aggregate too much
…: Helper compute servers have one round of communication to verify the range proof, then output the total. Original work didn't include DP noise, but can easily be added on
…: Helpers don't need to be online all the time — you can have an "ingestion server" that can't see any data, but can store records for computation later, and can do filtering on some plaintext criteria
…: Efficiency costs: Device is contributing to one bucket of a histogram. If we just used this naively, then PRIO would need us to send information proportional to the number of buckets in a histogram. But!...
…: This system has been used in a couple of places. Notably, Firefox's privacy-preserving telemetry, and Exposure Notifications Private Analytics from Apple+Google for COVID exposure notifications wth privacy. This included a way for health officials to get aggregate health metrics, running on iOS and Android devices, large-scale deployment of this technology.
…: If we want to compute huge histograms (and evidently ads use large histograms), how can we do it?
…: Distributed Point Functions, from paper Lightweight Techniques for Heavy Hitters
…: Can send a contribution to a bucket of a histogram while sending data only proportional to logarithm of number of buckets, rather than linear
…: Servers need to do more work to make up for it. Incremental DPFs are cool, go read the paper, it's fun
…: Each server can compute the aggregate for any particular index in this histogram, and they can evaluate the shares of data at that bucket. Can also do it for any subset of the buckets. To spare the non-cryptographers, I won't go into the details. We have an implementation, open-source, please go play with it.
…: Benchmarking: costs look within the realm of feasible. Depends on the size of the histogram. Exponential domain spaces require some sparsity of filled buckets, won't go into details.
…:Privacy Preserving Measurement at IETF: compatible with both Prio and the IDPF histograms I presented here.
…: There are always surprising costs, as we found out with Exposure Notification, that's why we've implemented it.

Nick Doty: Threat model for aggregator — first party that gets all the reports before sending them on. What can that single party learn? Are we trusting them to do anything re tampering or learning?
Mariana: Single party is an option, not necessary. It doesn't learn any content, but it does learn the timing of each data report, so goes into the decision about the browser delaying when reports get sent. Beyond that it sees only encrypted data.
Nick: Does it learn how much reports different browsers send? Traffic analysis?
Mariana: Can scope who are the dedicated parties. The parties might be ad tech themself. For exposure notification, the ingestion server is run by Google, so that they can do device attestation. Depends on who is looking at the aggregate.
Nick: They might learn something from what the device needs to attest, but none of the data that is destined for the people doing the computation.
Mariana: Shout-out to my Apple collaborator, who is also here observing!

Betul Durak: What are the experimental results covering? DP, sketching, etc?
Mariana: DPF and noise
Betul: The sketch — the range proof that it's not adding out-of-bound contributions. Adds a lot of complexity
Mariana: The proofs allow you to prove 1-hotness. We did not include them in the experiment. The more complex part is independent of the size of the histogram, and we were focused ont he question of how it scales up with histogram width
Betul: The sketch part is also quite expensive, so should also be evaluated
Mariana: The 1-hotness proof technique that I'm aware of doesn't include sketching. Let's follow up offline.

Kris Chapman: To add to Wendel's plea for thinking about business costs: I can't see how to explain this to advertisers, or to the general public. Are these just big-tech solutions that people won't understand? Do we just tell people "trust us"? Would love to see thoughts about how to convince advertisers or privacy concerns of individuals that this isn't doing things outside their expectations.
Mariana: Completely agree — explaining in a way that is easily understood and digested is hard! We succeeded at this with encryption, people understand that it protects your data without knowing how RSA works (OK maybe everyone does understand RSA but still). Protecting privacy of individuals includes a lot of things — MPC, Differential Privacy, other stuff too. We need a comprehensive view.
Charlie: What the ingester gets to learn, as Nick asked: If the client is sending a report to the ingestor that is a function of sensitive data — for instance, if there is only an attributed conversion, which is cross-site data — then the presence/absence of any report int he system could be sensitive, so in the ingestion model, the *count* of the number of encrypted records could reveal something about user behavior. We need to worry about this in general — MPC servers also can see the number of reports they are operating on. This is subtle, need to be careful about our designs, but we can mitigate enough so that the ingesters and the MPC servers really are learning minimal data. Just requires care.
Benjamin Case: Data on communication cost, as well as computation?
Mariana: Yes, will get it to you on chat / offline.
Aram: Would be very interested to hear studies on the success of communication success of the exposure notification during COVID
Alex Cone: I'll open an issue to crowd-source some literature (also dropped some stuff in IRC).

@betuldurak - Prio/DPF-related "Buckets"

Betul: Microsoft proposal for Oblivious Bucketization — similar to the PRIO buckets goal
…: We thought about this for the attribution reporting, like ARA (as Charlie has abbreviated Chrome's Aggregation Reporting API).
…: What a report means: each report consists of a key = a bit string, arbitrary length / arbitrary number of features, along with a list of values corresponding to that key. The values are what you want to aggregate — could be different types of values, will come back to it later.
…: Want privacy protection against aggregation. Each report will be secret-shared across two aggregators
…: In our protocol (different from Prio or IDPF) we introduce a third server, which makes things much more efficient. Browsers share all these reports directly to MPC, the reporting origin just interacts with the MPC provider.
…: We expect some malicious behavior from browsers, and a malicious browser can try to spoil the result of aggregation, need to detect and prevent such behaviors.
…: Also we have a threat model aligned with the previous ones: Aggregation servers are presumed to not collude with each other. More than that, we expect the possibility of a sort of malicious behavior from the servers — we want to differentiate between malicious behaviors of a server that deviates from the protocol in order to learn some information about users, vs. deviating in order to spoil the results. Motivation for server to spoil accuracy might be lower than motivation to spoil privacy, so we protect against the more dangerous one, don't worry as much about spoiling the total. We think easier than full protection against all malicious behaviors, but better than honest-but-curious.
…: Learn differentially private aggregates. Want to add noise in a controlled manner, in a "malicious-star" threat model, in that an MPC server might be colluding with the reporting origin. We would like to protect against reporting origin [scribe lost something here]
…: Flexible query structure: hierarchical querying mechanism, so that you can aggregate over individual features of keys, get histogram over them, and then aggregate over sub-buckets within that later. Flexible way of structuring, which I don't think other approaches can provide. Add noise at each hierarchical level of the aggregation; the noise added at a lower level propagates to the higher level later.
…: DP budget is computed in a very formal way, it's all in the paper. Same privacy guarantees as in Prio or IDPF (didn't know about Prio+DP noise, but Mariana says it's possible). More flexibility, to build hierarchy, what kind of aggregate on which feature of the key you get a histogram, get this on the fly from the ad tech vendor.
…: Obviously would like to have better efficiency, so want to use the least expensive computations. We did this by increasing the communication cost, but doesn't necessarily mean more expensive overall. Bonus: hierarchy allows better performance for very long keys.
…: Details analysis in the paper.
…: How it works, without details: Original reports go to server1 and server2, those two servers add some dummy reports which add masking to the true counts of aggregation. The reporting origin cannot tell which are original reports vs dummy reports, because of an oblivious shuffling (this is where the extra communications cost comes from). After we have these mixed and shuffled dummy reports, we combine the shares, bucketize, output results.
…: Two-party secret sharing, then integrate the third server to improve performance of the shuffling. Get secret shares, add dummy reports, shuffle via third server, reveal result from only two servers, show DP aggregates, and then stop or go to next layer in hierarchy.

…: We've tested the protocol in various settings. In sparse domain — where key size is large, but number of buckets that are used is small by comparison — tested 32-bit bucket labels but only 2^16 buckets used; slide contains data on performance on a particular system spec. Shuffling is the most expensive part.
…: Comparing to DPF via our friends from Google, single user contribution handling is 0.72sec (no sketching or DP) and bucketization with 1M records in 10 seconds.
…: Heavy Hitter with 256-bit key, 400K reports, zipfian distribution to reflect real-world data. IDPF takes 53min, our version takes 28min instead.
…: Overall comparisons slide of Prio-vs-IDPF-vs-Bucketization. We increased number of aggregators, used different techniques for robustness (no cost, unlike SNIPS or sketching), but there is a cost if we want sums instead of counts. Some additional benchmarks still in progress. Additional aggregator buys a lot of flexibility and efficiency.
…: On the Charlie left-right scale, this is a computation vs communication trade-off
…: DFP is a successor of Prio; DFP vs Bucketization are orthogonal. One focuses on a lot of computation on servers which is expensive in MPC land, bucketization uses a lot of communication instead. Server-to-server communication in PPML, but more communication does not necessarily mean more expensive protocols, as our work shows.

Charlie: Thanks for the presentation, very informative. Clarifying questions. For the benchmark with 2^20 possible keys: did that involve hierarchical evaluation, or just direct?
Betul: Just direct, no need for hierarchy at that scale.
Charlie: Threat model: We're adding these dummy records, so privacy is with respect to some privacy parameter, so MPC model is different, the end result learns something that is a function of the number of dummy records added.
Betul: Each server knows how much dummy it adds, communicates with the other servers about the dummies it adds in a masked way.
Charlie: One of the parties will learn the total number of (records + dummy records) in a bucket, so a value protected with DP?
Betul: Servers already learn the number of records but don't know how many for each bucket. Each server doesn't learn which are the dummy records from the other servers. Number of dummy records is by laplace mechanism.

EKR: Thank you, interesting. Case of correcting crash URLs / Poplar (new name for Heavy Hitters work): IDPF inspired, but we didn't contribute to the work, just suggested the use case! Don't want to grab any credit for work that wasn't ours. PRIO was invented by them first, then we suggested this problem, and then they solved this problem too.

Phillipp Schoppmann: Threat model: This requires more trust in computation server becuase of "honest majority" version, right?
Betul: Right, you need to trust two servers, not just one. Still protect against one malicious server.
Phillipp: Do you still provide security against one malicious server?
Betul: Yes, any one of the three servers (which play different roles, so it's not symmetric) might be malicious
Phillipp: What if malicious clients and server collude?
Betul: Some subset of clients. We don't think about a large subset of clients being malicious, though.

5m Break

Reminder to presenters: Please send slides to chairs!

Private Ad Measurements Constraints and Conditions discussion

- https://github.com/patcg/meetings/issues/17
- Discuss working group draft process.

Aram: intent is to step back and review holistically. Compare and contrast. Understand how to move forward around private measurement. Erik Anderson would like to present a comparison to start us off.
Eric: Joel will do this aim to be as neutral as possible - please flame if that's not the case
Joel: Not an expert, will present a distillation at a high level - problems that are trying to be solved, what was can do with them. What spaces they address
..: Features are contextual information. Labels are values. Three forms of helper: 1) trusted helper - complete trust to not reveal data or corrupt output 2) Semi-trusted helper - won't corrupt output, but might violate privacy (most proposals are here) 3) Malicious helper - messing with output.
..: will cover Browser event reporting APIs, Aggregate reporting and Meta/Mozilla IPA proposal
..: Apple/Google browser event reporting seem similar - brower-level mechanism to do attribution to click, with extensions for views. Metadata can be attached, stored in browser. When browser hits a destination site, triggered to report to downstream attribution collector.
..: Vary on number of source event ID bits. Could be used for training, generally can be used for aggregation
..: Trigger event bits vary. Random noise may be added. Most proposals are fire-and-forget. Timing delays used to address timing based attacks. Not all proposals have all features. E.g. some statistics might not be available
..: Browser aggregate reporting. Using helpers to enforce global privacy constraints. Most proposals have helpers only given a partial view. Trusted to not corrupt the output; this is validatable.
..: Trusted 3rd Party server e.g. using secure enclaves. Restricts allowable queries, adds noise. Simple to understand/implement. SPoF, and obvious attack surface. Perhaps this is a nice option to have available - e.g. fraud use cases, where qyuick response is important and we're willing to pay
..: MPC based proposals - uses secure MPC, returns a secret share. Allow side information. Each helper sees stuff that appears to be random noise. Distributor sees all data, but it's encrypted. Helpers are trusted to enforce privacy
..: Prio. adds SNIP - allows servers to verify that the sample is valid, protects against poisoning attacks
..: Bucketization - allows aggregation, in particular on subsets. Keys are subset of some key bits "bucket". Fake records inserted, shuffle and mask to protect privacy. Features/labels are hidden from helpers. Proven to work well with dense subsets. Works better with a third helper, although that could be chosen at evaluation time.
..: Incremental DPF - Core is a point function, can use clever math to create incremental evaluation over this function, and therefore benefit in sparse spaces (e.g. find 2^16 most common items in 2^128 space).

..: Masked LARk - presents an abstract, differentially private, Map-Reduce framework. Showed an example of training. Browser sends labels, including noise, server side combines. Features and models require trusted helpers. Labels, only semi-trusted

..: Private joins (Meta) - separates click from conversions in two differen apps, allows later joining using a match key, set per provider. Used to share information, but not readable. Presumably, user has to be signed on in both locations to set the keys. Both click and conversion data are sent separately, and joined in MPC. Blinding plus homomorphic encryption to construct the matches; allows aggregation. Requires semi-trusted helpers.

Aram: Thanks. Good summary. Next step is to discuss how these come together or exclude each other. 30 mins for this. Output should be opening some issues in proposals, so that we can advance a doc.

Martin: Shorter summary: we have a range of MPC options. Some of them have been demonstrated to work (for some value of that). Choices to make to move forward. Would like to start by addressing the question: is MPC the right thing to start with? Want to engage properly. Some of us have assumed it is, but don't want it to be an assumption. Answer depends on the requirements - no single entity that has access to a large slice of browsing history of multiple users, and advertisers get some useful information, and it works. I conclude MPC is the right answer, but let's discuss, before we move into details.

Aram: Topic: is MPC the way for continued progress. Do you have a technical objection to it? Or a clear objection beyond what's been stated so far?

James: Before answering the question, amazing to see the presentations, lots of clever stuff. There's an assumption in the space that can be covered in the question "have you stopped stealing sweets?" and therefore we need all these protections to stop folks stealing. I find that problematic. We're getting the cart in front of the horse - its' not to do with the math, we have a problem with information asymmetries between browser vendors and other players in the space, and so on

Brian: Seems to me that some sort of trusted server is going to have to be included. MPC sounds like it's going to be the way we go. Need to look at that closely in terms of unintended sideeffects. Are there going to be pockets of control in the ecosystem? Who manages these servers? What's the governance?

Charlie: I know that MPC is super appealing. I'm an MPC lover like Mariana. DOn't want to rush to conclusions. Suggest we spend some time looking at alternatives, and weighing them before we make a decision. Spent a bunch of time talking about MPC both days, to make a good decision, need to, as a group go through and reject alternatives. There are some aspects of MPC that we haven't figured out - will it work at scale, can it satisfy all of the use cases. We like it for the security mdoe, but does it check all the other boxes.

Aram: should open an issue in proposals to discuss this.

Erik A: was going to say something similar about explicitly rejecting proposals. Lots of use cases that we can't solve at event level with acceptable privacy. But maybe we need it for some things. Doesn't have to be a one way thing - not just yes/no - is it the right thing for certain parts of the problem

Andrew: Not are we choosing MPC, just is it a right place to start. I'm a fan of starting here. There are a few applications - gathering measurement, and then maskedLark allows machine learning on MPC. That gives us more of Charlie's checkboxes. I manage a DS team - important to be able to see data. Being able to specify a model mathematically isn't enough. There can be bugs in what we put into the MPC. Things like conversion prediction can be very sensitive to small changes inthe data, either a bug or otherwise. We might want to separate measurement from learning. For myself, would like to see how we can enable DS to play with data a little more freely.

Ben: Would love to make some progress. Would love to have some work product that we can collaboratively produce. So I'm hesitant with this idea of exhaustively traversing the space. That could take years. Chrome is

shooting for 2023 for cookie deprecation, which is tight. If we forgo the creation of standards until we've been exhaustive, that's going to take a long time. Maybe we can do things in parallel. If there's sufficient consensus around basic measurement with MPC, we could start on that, while evaluating other things in parallel. That's my plea.
Aram: Treat MPC for measurement and for ML as two threads we can handle in parallel. That's reasonable
Alex: +1 to Ben. Feels like on-device-only is being treated as innocent until proven guilty. MPC is being treated the opposite way around. Maybe because of market dynamics - ability of big players to move without standardization. Currently we're going with on-device without being critical. I don't want to be critical of on-device, but if we try to bottom out every angle in the MPC space (what if people create fraudulent reports, etc), that's too much.
Chris: Confused by this line of questioning, about whether MPC is the right approach. We haven't discussed the requirements. Martin gave a reasonable place to start earlier. Seems like it would be useful to nail down that, and have the solution fall out from that.
Reimundo: was going to say what Chris said. Can we reframe the question to, what is the security requirement for processing data off device? On top of that, re: how much time to invest in other solutions. We should set the requirements, and convince ourselves that MPC meets them. If we set up the requirements, then we can know whether we have a solution that meets them.
Wendell: Would love to hear more about the commercial viability of any of these solutions. Interested to work with the IPA crew to make something viable. Cost/longevity are key to these. Getting more details of that in whatever way makes sense - experiments…. More academic papers doesn't seem to be the right answer - doesn't help businesses to make operational decisions.
James: Want to point out that, there's not going to be just one solution - it's going to be up to market forces for what succeeds or fails. We will maybe have lots of solutions. Shouldn't be precluding that. To: what's the problem? Bunch of us have been working for years on that, that work got bogged down, but maybe we should revisit it, and reconsider the set of things that we could consider and advance. Lastly, andrew raised the point about how you know it's working, and advertisers ended up paying the wrong amounts, because of mistakes, and publishers lost out. Its not a simple question of picking one thing and moving on, there are other considerations.
Michael: Wendell said some of this: much as I'm an MPC fan - lots of clever crypto - I don't know what the cost multiplier for it is. If you want to do the same aggregation in MPC v.s. Single trusted server, what's the multiplier. A Charlie elephant arrow is: how much more are we willing to pay for X amount of privacy. We can do it in terms of $ for CPU/network. If the answer is 10x, then that would be a good thing to know, vs 10Mx, 1000X, etc. Sliding scale. The reason I'm not entirely convinced is that we don't know how many extra 0s there are. Need to know that.
Mariana: Many people asking about cost, and business needs. Perhaps set up experimental frameworks to look across companies and evaluate. Agree that papers can only go so far. We tried to get closer to evaluation, but all of us thinking together about what we want to know would be a good way to get answers to these questions
Nick: Martin articulated one privacy goal (single server with history over many users). That's a valuable goal. There might be others. Data minimization wouldn't be the only one. That's why we're going to have this whole deliverable about that. Since we have that one articulated already, does anyone think that should be a non-goal. That would be useful to know. Or if you have other privacy goals to share - user understanding, or user control would be maybe good goals. Even if we're going to work on things in parallel, it would be good to flesh out those privacy goals. I'm happy to work on that.

Martin: Nicks' request is reasonable. I think that requirement I articulated is the overriding one from my perspective. As we go through this, we'll find some at lower levels, but I don't think we'll find another one that's as important as this. There are other things we have to work through, e.g. if it's 100Mx as expensive, maybe that's unacceptable, and we have to factor that in, but we won't really know until we start doing something. Same with the trying-everything. This is not where we want to go, I suspect. Would rather make the decision to start with something, and consider alternatives in parallel. If someone comes up with something better, I'm good to change.

Aram: instead of break, can we power through? Good. let's go. 3 proposals in teh space. PLease open up issues if you're an owner. MPC for counting/ML. Other questions raised. Let's the get the issues up and start the discussion there.
..: Topics API proposal has been brought to teh group. Yesterday, few folks said we should focus on metrics before anything else. Sounds like rough consensus to not pick up Topics for teh next meeting at least. Any objections to deferring that until after the next meeting at the earliest. No Objections.
Bmay: do we have to decide that now?
Aram: no. but that'll be the default position until there's a substantial issue in the repo
Brad: timing matters - if it's not going to be picked up, it should go to another group
Aram; lets aim for discussion on an issue over the next week and establish whether we want to pick it up
EKR: let's be clear about "discussed" v.s. "Taken up". Do they come to meetings and get meeting time. My concern is available bandwidth. Doesn't really matter which forum its in. Can we really make progress on both things at once? This meeting has been productive, but wouldn't want to slow things still further by giving everything 20 minutes.
Brad: major concern is getting it adopted by some group. Meeting time is a second concern. If folks want to have side calls e.g. in WICG (we haven't covered that as a process item for this group), we'd be happy with that
Aram: we've allotted a label and a process for managing sub-groups, especially since we've agreed on infrequent meetings.
James: if you have 6 things, and one unit of capacity, it's hard to engage. The more things that go on in parallel, the harder it is to get genuine consensus. Becomes a tax on smaller organization to have many things in parallel. Want to balance that with innovation, need to be mindful not to exclude folks
Martin: Folks used the word adoption - it's loaded with meaning. Adoption to me means some formal acknowledgement from the group that this is "good". Cf ITF. I don't want to see anything this group adopts falling short of that. The general shape of the solution has to be right. That's different from what Google wants, which is a place to discuss this that has IPR protections. We can set that up without those things being formally "adopted" or blessed by this group.
Aram: yes. In parallel
Brian: Seems that lots of people are working in parallel directions. We should adopt something we like and find a way of working together. What's our cadence
Aram: will come back to that
Brad: To clarify "adoption" - as per the charter: something the group wants to discuss, not an endorsement. A thing that is to be discussed with IPR protections with at least one other interested party.
EKR: This CG can attempt to progress an idea towards adoption. Don't care what the word is. But we need a word for "work item for the group", not just a thing being talked about. That's an important distinction, and we need to be clear about it. Words should be clear.

Sean: OUtput of a CG is a report.
EKR*:* Basis of a report requires consensus
Aram: Adopted can mean having a home for discussion, not recommended
[EKR: This is not at the mic, but the text of the charter is quite clear that adoption means that there is consensus “that the work should be taken on and that the document is a good basis for a Community Group Report”, so it’s a higher bar than “there is agreement that this deserves some airtime so it gets IPR protection].

Group Report.

..: Next meeting - we’ll open an issue with polling. Back and forth on whether this should be monthly or every two months. Wondering if, given where we are, we should bias to monthly. Is there an objection to every month?
Martin: 6-8 weeks
Bmay: while there’s stuff to talk about, we should push. CGs can take a long time to get stuff done.
Aram: 6 weeks… any objections?
Sean: Suggestion to splitting by a day. Will put the doodle poll up. Promise we’ll move the times around
Erik: assumption is that we’ll put issues in GH, and have meaningful discussions there. Imagine that there are some small groups who might benefit from breakout live sessions. Are we okay with that? Discouraging? No position?
Aram: same as privacy CG. There is a label. We’ll set up an issue template for an ad hoc meeting, can be discussed in that thread.
Bmay: Concern - how to know that they're happening
Aram: issues will email the list. When theres’ an ad hoc meeting, lets push a pull request to teh meetings repo
Erik: sounds fine as a place to start
Aram: that’s satisfactory. Sean to close
Sean: thanks again. Looking forward to progress. Closed. Bye. Thanks all.

5m break

Scheduling next meeting and discussing focus

- Potential issues to pick up:
    - https://github.com/patcg/meetings/issues/19

Queue

Queue Closed

Participants Day 2

1. Aram Zucker-Scharff (The Washington Post / Co-Chair)
2. Sean Turner (sn3rd)
3. Brian May (dstillery)
4. Sean Bedford (Meta)
5. Chris Wood (Cloudflare)
6. Robert Kubis (Google)
7. Matthew Liu (The Washington Post)
8. Alex Cone (IAB Tech Lab)
9. Andrew Pascoe (NextRoll)
10. Sam Weiler (W3C/MIT)
11. Charlie Harrison (Google Chrome)
12. Simon Harris (DPG Media)
13. Lisa Markou (FordDirect)
14. Sergey Tumakha (Microsoft Ads)
15. Alex Brasil (Shopify)
16. Russell Stringham (Adobe)
17. Valentino Volonghi (NextRoll)
18. Tara Whalen (Cloudflare)
19. Sanketh Menda (N/A)
20. Don Marti (CafeMedia)
21. Martin Thomson (Mozilla)
22. Denis Charles (Microsoft)
23. Ratko Vidakovic (AdProfs)
24. Nazar Andrienko (Twitter)
25. Mark Nottingham (Fastly - observer)
26. Kris Chapman (Salesforce)
27. Shivan Sahib (Brave)
28. Michael Kleber (Google Chrome)ƒ
29. Mark Kent (Twitter)
30. John Delaney (Google Chrome)
31. Alex Turner (Google Chrome)
32. Wendell Baker (Yahoo)
33. Michal Bryc (Twitter)
34. James Aylett (Omnicom Media Group)
35. Phillipp Schoppmann (Google)
36. Mariana Raykova (Google)
37. Brandon Maslen (Microsoft Edge)

38. Jacob Naim (Demandbase)
39. Chris Needham (BBC)
40. Lorenzo Hernandez (NextRoll)
41. Brad Lassey (Google Chrome)
42. Akshaya Mani (Optable Inc.)
43. Przemyslaw Iwanczak (RTB House)
44. Benjamin Dick (IAB Tech Lab)
45. Christine Runnegar (PING co-chair) (part meeting)
46. Leon Yin (Microsoft/LinkedIn)
47. Wendy Seltzer (W3C)
48. Jeffrey Yasskin (Google Chrome)
49. David Dabbs (Epsilon)
50. Bryan Kapicka (Twitter)
51. Erik Anderson (Microsoft Edge)
52. Bosko Milekic (Optable)
53. Joel Pfeiffer (Microsoft Advertising)
54. James Rosewell (51Degrees)
55. Christy Harris (Future of Privacy Forum)
56. Benjamin Case (Meta)
57. Anuvrat Singh (Amazon)
58. Eric Bruce (NYT)
59. Betul Durak (Microsoft)
60. Steven Englehardt (DuckDuckGo)
61. Nicolas Arciniega (Microsoft Edge)
62. Kelda Anderson (Microsoft Edge)
63. Davis Gilton (Microsoft Advertising)
64. Nick Doty (Center for Democracy & Technology)
65. Josh Karlin (Google)
66. Ericka Wright (Intuit)
67. Aditya Desai (Amazon)
68. Ben Savage (Meta)
69. Moshe Lehrer (Neustar)
70. Qun Wei (Lemonade)
71. Robin Berjon (The New York Times)
72. Narayanan N (Huawei)
73. John Mooring (Microsoft Advertising)
74. Bill Densmore (ITEGA.org)
75. Eli Grey (Transcend)
76. Richa Jain (Meta)

Participants Day 1

1. Aram Zucker-Scharff (The Washington Post / Co-Chair)

2. Sean Turner (sn3rd / Co-Chair)
3. Ben Savage (Meta)
4. Valentino Volonghi (NextRoll)
5. Chris Wood (Cloudflare)
6. Mark Nottingham (Fastly - observer)
7. Andrew Pascoe (NextRoll)
8. Tara Whalen (Cloudflare)
9. Wendy Seltzer (W3C)
10. Sergey Tumakha (Microsoft Ads)
11. Russell Stringham (Adobe)
12. Martin Thomson (Mozilla)
13. Michael Kleber (Google Chrome)
14. Denis Charles (Microsoft Ads)
15. Brian May (dstillery)
16. Sean Bedford (Meta)
17. Brad Rodriguez (Magnite)
18. Don Marti (CafeMedia)
19. Shivan Sahib (Brave Software)
20. Michal Bryc (Twitter)
21. Lukasz Wlodarczyk (RTB House)
22. James Rosewell (51Degrees)
23. Phillipp Schoppmann (Google)
24. George London (Upwave)
25. Lisa Markou (FordDirect)
26. Akshaya Mani (Optable Inc.)
27. Matthew Liu (The Washington Post)
28. Erik Anderson (Microsoft Edge)
29. Anuvrat Singh (Amazon)
30. Joel Pfeiffer (Microsoft Advertising)
31. Robin Berjon (The New York Times)
32. Betul Durak (Microsoft)
33. Brad Lassey (Google Chrome)
34. Nick Doty (Center for Democracy & Technology)
35. Nicole Lesko (Dotdash Meredith)
36. Jeffrey Yasskin (Google Chrome)
37. Wendell Baker (Yahoo)
38. Alex Brasil (Shopify)
39. Nazar Andrienko (Twitter)
40. Mark Kent (Twitter)
41. Bryan Kapicka (Twitter)
42. Bosko Milekic (Optable)
43. Qun Wei (Lemonade)
44. Eric Rescorla (Mozilla)
45. Alex Cone (IAB Tech Lab)

46. Mariana Raykova (Google)
47. Sanketh Menda (N/A)
48. Christine Runnegar (PING co-chair)
49. Benjamin Dick (IAB Tech Lab)
50. Leon Yin (Microsoft/LinkedIn)
51. John Delaney (Google Chrome)
52. Anthony Molinaro (OpenX)
53. Nicolas Arciniega (Microsoft Edge)
54. Alex Turner (Google Chrome)
55. Harneet Sidhana (Microsoft Edge)
56. Lorenzo Hernandez (NextRoll)
57. Kelda Anderson (Microsoft Edge)
58. Simon Harris (DPG Media)
59. David Dabbs (Epsilon)
60. Shubho Sengupta (Meta)
61. John Douglas (WeTransfer)
62. Raja Johns (FordDirect)
63. Benjamin Case (Meta)
64. John Wilander (Apple WebKit, here as observer)
65. Alissa Cooper (Cisco)
66. Steven Englehardt (DuckDuckGo)
67. Sam Weiler (W3C/MIT)
68. James Aylett (Omnicom Media Group)
69. Wayne Blodwell (TPA Digital)
70. Jacob Naim (Demandbase)
71. Davis Gilton (Microsoft Ads)
72. Aditya Desai (Amazon Ads)
73. Garrett Johnson (Boston University)
74. Ericka Wright (Intuit, representing self)
75. Patrick Jordan (Microsoft Advertising)
76. Julien Delhommeau (Xandr)
77. Moshe Lehrer (Neustar)
78. John Sabella (PubMatic)
79. Eli Grey (Transcend)
80. Alex Austin (Branch)
81. Christina Ilvento (Google Chrome)
82. Narayanan N (Huawei)
83. John Mooring (Microsoft Advertising)
84. Eric Bruce (NYT)
85. Charlie Harrison (Google Chrome)
86. Ratko Vidakovic (AdProfs)
87. Stephen Somerville (MeritB2B)
88. Daniel Gale (Simplifi)
89. Eric Rescorla (Mozilla)

Cursor Nature Reserve:

🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳

🌳   🐇   🌻   🦖   🌳

🌳 🐄   🦄   🌴   🦖   🐝   🌳

🌳 🌲   🐌   🌳

🌳   🐸   🐓🐔   🌷   🌳

🌳 🐕   🦖 🐐🐑   🐒   🦢   🌳

🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳🌳