Anti-Fraud CG Meeting

- Chairs: **Samuel Schlesinger**, Sofia Celi
- Scribes: Manu Sporny

Agenda

- AFCG Focus Areas into 2026
- Unconference
- Next Steps

Queue - add yourself at the bottom

-
- Vinod
- Manu - what is the lowest command denominator solution?

Attendees (sign yourself in):

1. Samuel Schlesinger (Google Chrome)
2. Eric Trouton (Google Chrome)
3. Michael Ficarra (F5)
4. Manu Sporny (Digital Bazaar)
5. Philipp Pfeiffenberger (Google Chrome)
6. Vinod Panicker (Amazon Ads)
7. Kosei Akama (NICT/Keio University)
8. Tara Whalen (W3C)
9. Bartosz Kwitkowski (PREBYTES)
10. Yosuke Ito（OP-CIP/MRI）
11. Hiroyuki Sano (Sony)
12. Joel Antoci (Shopify)
13. Nick Doty (CDT)
14. Isaac Z Foster (MSFT Ads)
15. Lionel Basdevant (Criteo)
16. Maxime Guerreiro (Cloudflare)
17. Mark Nottingham (Cloudflare)
18. Yumi Dobashi(OP-CIP/MRI)
19. Tomoya Asai (WebDINO)
20. Tatsuya Katsuhara (Digital Agency, gov of Japan)

21. Alan Buxey (MyUNiDAYS)
22. Ben Kelly (Meta)
23. Ben VanderSloot (Mozilla, observing)

Scribe:

Sam: Thank you for coming here in person to anti fraud CG. We'll be talking about Antifraud today. Please sign in. Agenda is on Github and will be in meeting minutes. These are public meeting minutes, please make sure not to say something you are not ok with being public.

Sam: We are going to talk about current focus areas, do an unconference, and then talk about next steps. Who is going to do what, how to manage participation going forward?

Sam: Our charter, identify and design scenarios to figure out unwanted scenarios. Focus on security privacy and accessibility. We want to get stake in fraud and abuse (preventing it), protect sites and the ecosystem in whole.

Sam: In 2026 we met in May – three high level themes; impact of AI agents on fraud, items with clear user benefits (account security, payment, challenge bypass), ecosystem benefits at scale (ad fraud, bulk account creation, etc). Other areas, incubation o cross-platform for rate limiting, residential proxies. Open floor on this slide – what should we prioritize - dangers here, what should be discussed? Accurately capture group intent and interest.

Manu: Where do digital IDs, verifiable credentials, wallets, etc come in?

Sam: Rate limiting - rate limit against high quality resources, fraud in ecosystems…. Those are in scope.

Nick: Nick Doty, Center for Democracy and Technology – methods for server to declare type of anti=fraud, APis for purposes, some potential to signal for why you are doing something? Might be quite different from AI / rate limiting. Use cases in past were limit use cases that could be abused (ad tracking and targeting) – "I promise I'm just collecting data for certain use"

Sam: Sympathetic to this, very fraught - two paths – user agents try to force declarations as accurate… some feel they can't do that for some reason or another – both paths have issues. Hard to imagine we do useful work there – don't see path through. Don't want to see declarations and then UAs can't do anything about it.

Nick: If we do a lot of work and it doesn't improve things, we are already in the degenerate state – this is going to be hard work.

Eric: Hard from a browser enforcement, if there is a trusted 3rd party agreed on by standards body to do enforcement would be easier to implement – who i the right body to do enforcement.

Nick: Yes, individual browser vendors, browser software won't do enforcement - we could use out of band enforcement mechanisms. We can have some sort of consumer protection – yes, hard conversations, but interested in having those discussions.

Sam: Yes, if we can get another entity to enforce, we should discuss.

Michael: Continuing from what Nick said, I'd also like to explore additional capabilities that we could make available to trusted security and antifraud providers. There's obviously a political issue with determining who those providers would be. In an ideal world, we would have intentional powerful APIs instead of misusing accidentally powerful APIs. What would those look like?

Sam: Exposure of novel signals on the Web.

Michael: To me, signals implies misusing existing APIs that are intended for something else. I'm talking about something powerful and intended for this purpose.

Sam: So capabilities?

Michael: Yes.

Nick: Could be worse than status quo. Makes me a bit more cautious..

Isaac: Microsoft on ad side, publicly declaring something, with privacy sandbox – public attestation we take seriously – worth mentioning we took it seriously (public attestations).

Eric: Would like to figure out incubation for proposals. What do the new ones look like.

Bartosz: A few introductions. Interested in antifraud related to institutions. Remote access account – involving fingerprints, residential proxies, remote access is unique for the approach. As far as handling this, on various methods, there is no agenda for this. Thinking around attestations for browsers, focusing on browsers as input for interfaces where the user can prove itself as not being spoofed in any way. Some overlap on topics, from financial institutions point of view, space for antifraud agenda.

Isaac: Involving TPMs on consumer devices, also interested in server side. +1 on implementation detail and not necessarily a capability.

Sam: We have a history of discussing environment integrity signals in this group, we can summarize. Could discuss/figure it out – lot of interesting work, fraught space.

Sam: We are going to go through a list of topics – brief topic, deliberate, pick rank order and go from there.

Vinod: Wondering whether we need to think about coarser grained entities in flow of money vs. anything else. Interested in how credentials, digital credentials, can be used for enforcement on potential threat actors not end users.

Sam: Under digital wallet? Credentials for intermediaries

Vinod: In digital ads we have multiple intermediaries in the supply chain, flow of money, but not end users themselves. These entities in the flow have incentive to inflate traffic to get ad revenue – enforcement actions on these entities end up side stepping concern around user identity. Entities of transacting parties, less heightened concern around privacy in general, taking actions will give outcomes we're seeing to stop fraud. End up getting same outcome.

Sam: Alternative threat model – ecosystem actors – threat modelling entities that are not users – or do we want to say intermediaries like ad tech?

Vinod: This model could be applicable outside ad tech – what about other dimensions – intermediary credentials v. user-level credentials

Eric: What about the rate limiting proposal – how do we build guardrails for the ecosystem – user opt in, opt out – can new startups get access to these signals. How do we make sure people don't get locked out of devices, etc?

Sam: Discussions about open web and future of open web – how can we make web more open wrt. Fraud?

Manu: How focused on tech will this group be?

Sam: Some are familiar w/ tech details, others are not – we can discuss

Mark: Been following work from afar - wondering if there has been much discussion of differentiation of use cases. Some folks might be here for transaction fraud, but that seems like it might be qualitatively different from ad fraud. Some sites might also use tech here to solve other stuff, we might not want to assume whatever site wants to do as legitimate. How can this be defined in different scenarios?

Sam: Declarations of intent?

Mark: Yes, possibly. Let's not necessarily put them into the same bucket?

Sam: Different streams for different targets?

Mark: We could talk about WebbotAuthn – can you separate these things… not so sure.

Eric: Browser based solutions – oversimplification, would like to use work that privacy pass is coming up with – and to make sure we're working w/ other orgs, some of us are in multiple groups w/o being duplicative –

Sam: charter permits us to use WICG to incubate, but if we want to standardize, we might need to become WG.

Eric: Nothing about AI in there right now.

Sam: Focus has changed, but don't know if the core problem has changed.

Sam: Anything else pressing for anyone? Let's discuss what we should talk about

Manu: What is the mandate?

Sam: We're trying to bring together the right people – we could have the most impact w/ browser API. We should be moving in the direction of WebAPIs.

Sam: Building guardrails for ecosystem solutions has been the main stumbling block historically? (Sam thinking out loud about focus)

Philipp: AI warrants its own reputation -

Sam: Private rate limiting for AI agents.

Benjamin: Don't know what state web-bot-auth is in – very much talking about agent but no real discussion about how to tie in user credential its acting upon. Agent acting on its own is a crawler, but acting on behalf of the user is its own situation. Maybe this is all server-side.

Sam: That has been discussed ; on our mind. We could have impact on APIs, automate browsers. We have a risk of people using Webotauth in a way that is not good.

Philipp: Guardrails for ecosystem - values to uphold.

Sam: We should spend a lot of time on that – need strong group consensus. Let's start with some shorter topics. Webbotauth per user gap.

Eric: One thing if we were to implement Webbotauth on browser, extensions, multiple browser extensions, HTTP authenticate – page interactions – might not be relevant there. Local agents get blocked by sites, signals behind user agent, ignore automation or be comfortable with automation – some things are ok, some things are not – does seem like local agents operate differently (browsing on behalf of user vs. browsing independently). Differentiating between

automation for good vs. potential automation that's bad. Attesting to an agent might not be a good solution, ecosystem openness concerns. Different ways to approach this would be good – something we need to solve.

Sam: Webbotauth is being generated in IETF – Cloudflare/Google – bots and crawlers can declare identity and get differentiate level of access from website, get more access – have to block bots or get overrun but there are legitimate uses.

Phillipp: On local browser, can reattach stuff, not clear which request comes from agent vs. something else – early days, front-loading concerns might be helpful

Eric: I think we should let the IETF deal w/ remote agents.

Vinod: Local vs. remote, autonomy vs. automation

Phillip: If it's some process on my machine, from defensibility, disambiguate between agents running in browser vs. agent running in data center.

Vinod: You're looking at relying party of origin, in reality of access – does it change anything from users perspective – I can use different technologies to achieve certain things. From a fraud perspective, is that key concern from the users perspective?

Sam: There is a threat modelling issue – no reason for RP to trust that local agent to trust operation of agent – you might trust agent provider, might not trust user but do trust agent provider. Both are relevant for the web, both use browsers,

Ben Kelly: You need to differentiate between user vs. agent.

Vinod: How is browser being operated

Ben: Not sure about browser attesting, could prevent new browsers from being created – its a risk

Vinod: Signal being available vs. signal being used.

Phillipp: I get the concern, don't want to get to design.

Eric: Building guardrails for ecosystem solution s– we don't want to prevent new systems.

Mark: Webbotauth first WG meeting was last week – there was pushback on some items. In an uncertain state on where its going to go. Don't presume it will look like what went out – might be based on HTTP Signatures, we'll see – it's bot focused not browser focused. Side meeting for A2A protocols, strongly attended 150-200 people, lots of interest in doing Oauth for agents –

wary of this group to define authentication mechanisms for agents – strong push for things like that there.

Nick: Larger fears is reputation signing w keys – current draft says MUST NOT use user key – concerned about how to apply to individual browser traffic. End of the slippery slope.

Mark: Already happening.

Sam: Per user gap is definitely not what we want to do.

Sam: What is scoping of the group – gist is "We will not develop cryptography here, discuss at a high level here".

Manu: We should be careful about where "browser/app-layer" crypto happens (W3C) vs. where cryptographic primitives happen.

Sam: yes, we should get input from various communities.

Sam: Moving on – large set of use cases – is ther ar use case that isn't covered?

Bartosz: Sharing users screen or input – convince user to use credentials – banking/crypto – instead AI takes over computer screen – browser used for fraud afterwards – sounds simple, but difficult to approach for most users – before covid, remote software/servicing wasn't difficult to differentiate wrt. Fraud… during COVID, explosion of remote software – leaving computer in office and performing actions. When user is using remote computer at office for legitimate transactions – when computer is authenticated – when user i moving w/ laptop – home laptop hijacked by criminals.

Mark: To turn it around a bit, current use cases – from perspective of provider of service – incoming traffic from browser – if you bucket them, if its online/payment interacting with bank account – need high level of assurance. Strong authentication is required for situation. Then you have advertising based used cases, different relationship between user and service. They don't have first party relationship, don' have aligned incentives. Then you have account creation - brick and mortar relationship. Then you have random websites – DOS is one of the issues – some random website having access to same facilities and information – that doesn't seem to be justified. Not saying we can do anything about this, but worth talking about.

Michael: We have to consider vital services are not doing these anti-frad measures themselves, few institutions doing anti-fraud. More tenable to enumerate all parties that legitimately need.

Mark: Depends on how services are composed – might be first party or third party services used.

Nick: I had a similar classification in mind – aggregate measures vs. high value individual transactions – different protection methods, but align w/ user interests. When someone spends lots of money, alignment w/ bank when doing large spending. However, with ads, it's a bit different – that distinction might be useful – user cares more about one of these than the other. User is engaged in high value transactions. Declarations of intent, promises, what you'r willing to commit to. Aggregate use cases vs. individual high value transactions are significant wrt. Promises to user/regulator.

Bartosz: Look at this as more general; a way to prove true location of user – states where government is allowed, provider allowed, to apply to users in a given state. In general, browser providing server flag that might be considered as remote.

Sam: Age verification – verified location, verified age, privacy concerns there.

Manu: what's the minimum viable use case for a working group charter? fraud protection use case + not violating open web practices. but those conversations seem wildly different.

Sam: This group is trying to figure out a way to bridge that gap

Sam: Guard rails for ecosystem solutions – wanting to attest to integrity of browser – strong interest by website involved. Google had proposed Web environment integrity and was shelved. Adjacent to this, we shouldn't try again for something that has that shape – probably wouldn't work. Perhaps one guardrail – we don' specify subset of acceptable hardware where APIs can run. Should work on any hardware.

Michael: Makes me uncomfortable – when considering an API, there will be some upside to it – we need to make that decision later.

Phillipp: We need to be specific on hardware – abusing API to limit access to particular platforms – do I feel like population is able to do it – design APIs that

Eric: we had agreement in anti-fraud space – be mindful of full set of representation across all groups. Getting signals from folks early on in design process. We might have controversy – we might have malicious behavior that needs to be blocked, point of doing detection is to know when to take action – we will have to figure out what the tradeoff are.

Michael: How do we get input from the WEI opponents?

Phillipp: Invite these folks into discuss with us.

Nick: I did document concerns (which I tried to do in a more polite, collaborative and friendly way), wasn't necessarily heard until it was said more strongly.

Vinod: Some guardrails are difficult as stated – how can we frame this in a way that is more people-centric

BenK: Getting feedback from the community is important – some feedback was more kneejerk – get threat model first – what are goals/objectives. There are tradeoffs – how are we going to handle tradeoffs. You could get feedback – get principles – eg: maximize privacy but allow people to do what they're doing.

Manu: priority of constituencies should help resolve those:
https://www.w3.org/TR/design-principles/#priority-of-constituencies
https://datatracker.ietf.org/doc/html/rfc8890

BenK: Priority of constituencies – things are not solved by priority of constituencies.

Sam: to Marks' point, differentiating use cases – where is extracting value happening? Who is doing it? I fear we'll talk about it and not refine it.

Phillipp: We don't want to build something where Google/Amazon/Apple say "YOu can use this stuff only on these devices".

Mark: That's a good place to start – web is an open system, new actors can come in – only good people use the APIs is not something you want to say.

Manu: concept of amount of friction that APIs cause. high-value transactions should be high friction (like a $10,000 transaction), should always be high-friction and never automated or automatable. more privacy-preserving,
Sam: mostly we are trying to reduce friction in many applications
Manu: but that's dangerous/concerning for high assurance use cases, when low friction. don't want the user to automatically hand over their national ID card.

Vinod: Is that always going to be true – could it be provided in an alternate way.
Manu: in deployments, already hearing from implementers that they would rather just ask for a government-issued ID on every transaction, rather than figuring out something more subtle. nightmare scenario, and hard to roll back.

BenK: High friction use case, perception – incentive to have higher friction – should this group be defining friction level?

Sam: Trust websites to create their software for the level of friction that works for them – reducing friction for people on the Web should be our goal.

Nick: Increasing friction should be in menu – not always on API design to reduce it – relative dimension. Device attestation, came up with questions – happy to pick that back up or work with

others – identify when normative principles might be crossed. Principles – friction would be good to add to the list – should this be used frequently – new for web platform to talk about how to use things infrequently – designing to be used as many times as possible – infrequency adds to societal value. Government ID work is where that's coming up a lot. Don't want age verification to be used constantly – age verification system that is rarely used.

Eric: Some of the things we build – prove reputability, trusted in some way – for everyone else w/o access to signals - maybe they need to fill out captchas – curious if we consider this open/fair users on – worse experience, is that fair?

Nick: Rigidity was one of the considerations – are people being punished – level of punishment – those users get denied access? Or is it slower experience – level of rigidity, design something where it's not relied up on in rigid way is important difference for societal impact. How does lockout happen.

Vinod: A set of constraints from within we must operate will help narrow down use cases.

Sam: That brings us into our next section. Who can take action items? Nick, are you interested in taking point on this?

Nick: I have an initial draft, could collaborate but can't lead/work solo, because the group would end up stuck waiting for me.
https://gitlab.com/pitg/device-attestation-user-considerations/-/blob/main/device-attestation-user-considerations.md

Manu: willing to liaise with digital credentials for high impact high assurance use cases

Sam: Anyone working on AI generated fraud on their backends?

Phillip: Automated fraud using LLMs –

Vinod: I can contribute from an ads perspective.

Nick: Would be useful to hear Chromes perspective on this. Willing to talk to Chrome folks and Isaac from Microsoft – could see if there is something to present.

Eric: The browser alone cannot be the one to mediate these decisions – from a competition perspective, too contentious. Has to be agreed to by standards across multiple browsers and/ 3rd party - did not want to go in this direction. Happy to discuss more or make introductions to the right people.

Michael: Michael, exposure of novel – thought about TPMs and how they could be used on the Web. Haven't considered non-TPM-based capabilities? Is that out of scope – interested in conversation – don't know where to start.

Phillipp: What are the things we want the TPM to do?

Michael: Rate limiting, yes – could work backward toward other ways to make that achieve that. Been doing work on residential proxy work as well.

BenK: Please link to them in the notes.

Sam: Thanks for everyone's participation!