Notes - Fenced Frame API (TPAC breakout)

Dom Farolino, Google: Want to make the community aware of fenced frames, API shape. Some of the talk will be a delta from previous public talks. Also technical challenges like permissions policy, intersection observer. Want feedback from engineers, partners.

Shivani Sharma, Google: In a world without 3P cookies, everything partitioned by site, but some use cases require displaying content from different partitions, including cross-site data. Initially motivated by interest group based advertising (FLEDGE, similarly PARAKEET). Shared Storage has an API, selectURL, which allows a worklet to use cross-site data but then hide the result by displaying in a fenced frame.
...: Origin trial is ongoing, can be used with FLEDGE.
...: Next step, FLEDGE requires fenced frames.
...: TAG review is satisfied, with a request to gather feedback from other vendors, resolve some outstanding issues, and file another review.

Michael Kleber, Google: Historical footnote: first proposal predates FLEDGE or its predecessors. Originally Ben Savage from Meta suggested something similar inside the Web Advertising Business Group.

Shivani: here is an API sample for interest group based advertising
...: embedding page calls navigator.runAdAuction, uses worklets to run auction, browser returns a urn:uuid: which is opaque to the embedding page
...: embedding page can create a fenced frame and pass it the urn:uuid: source, which the browser maps to the ad URL
...: the URL chosen is k-anonymous; the fenced frame doesn't know who the user is on the embedding page
...: no communication from the fenced frame to the embedding page, so cannot share the interest group back
...: fenced frame has network access in this example

Shivani: network access
...: server-side identity joining (timing/IP side channel) is a concern, which is an ongoing challenge

...: network is not restricted right now, but long term we want a solution (e.g. web bundles, trusted network solution)

Christian Biesinger, Google: in your example, fenced frame had a urn:uuid: source URL; is that the only kind of URL that works in a fenced frame?
Dom: you can use a urn:, https:, about:blank or localhost

?: what is a trusted network
Shivani: TBD, but the idea is a network/server that we can trust not to log

Simeon Vincent, Google: Can origins be expanded to include extension origins?
Dom: I don't think we want more capabilities than iframes have
Shivani: note that any ad in a fenced frame could still be blocked by extensions

?: applications list is ad-focused; how many use cases outside ads have you explored? You might need a result code (e.g. payment went through) in some cases
Shivani: future slide may answer

Shivani: speculative alternative information flow
...: want to display some unpartitioned data to display e.g., last 4 digits of credit card number
...: embedding page can send contextual information to fenced frame when it doesn't have user-specific data, to fetch resources etc
...: then later, it will have no network access or write access to storage
...: this is an example where the URL need not be opaque; if a fenced frame gets access to cross-site data we want to prevent exfiltration

Shivani: isolation from embedding context
...: storage isolation: origin-based storage/communication (e.g., BroadcastChannel) cannot reach from the fenced frame to frames outside it, even if same-origin
...: use storage partitioning with a unique nonce to the fenced frame tree, not shared
...: frame tree is isolated; window.top etc do not reach outside, cannot postMessage outside
...: to avoid privacy leak, things like CSP: frame-ancestors, cannot honor it to the top of the page, only to the root of the fenced frame tree
...: along with other behavior differences, this motivated us to require an explicit opt-in response header (Supports-Loading-Mode: fenced-frame)

...: actively considering intersection observer; embedding page programmatically moving the fenced frame could be used to signal information
...: are asking ad tech companies to learn more about the use cases
...: slide links to a document which explores the solution space, such as gating events on whether the move is programmatic
...: aggregate reports are another solution, but breaks use for clickjacking protection
...: would love to hear from folks about things that would break once IntersectionObserver is not available; right now it behaves the same as iframes

Shivani: Permissions
...: delegating permissions to a subframe could be used as a communication channel
...: initially wanted no permission-based features
...: if embedding frame is fine with iframes using attribution reporting, we allow it in fenced frames
...: don't even load the fenced frame otherwise
...: going forward we need to make it more flexible, allow embedder to use allow="" attribute similar to iframe
...: whichever consumer API wants permissions enabled in a fenced frame, they can include it in the k-anonymity check (with the URL); then fenced frame gets the permissions but doesn't send contextual information
...: also introduce an allow attribute which confines whether content can load

Brian May: What does the k-anonymity checker work?
Dom: per-URN-generator concern, specific to FLEDGE/PARAKEET/etc
Michael Kleber: nothing to add
Brian: is the check server-side?
Michael: some server-side support is needed to determine this URL is seen by enough different browsers, without leaking information
Brian: specification?
Michael: will find link
Brian: some number of calls needed to server-side infrastructure; is there some validation that it's not the same client calling repeatedly?
Michael: yes, that too. Read more here: https://github.com/WICG/turtledove/blob/main/FLEDGE_k_anonymity_server.md

Dom: configs more generally

...: today you operate a fenced frame by obtaining a URN from FLEDGE/PARAKEET/Shared Storage/etc
...: generated with cross-site data, embedder doesn't have access
...: handled through URN mapping to correct URL
...: want to also store permissions policy, width, height, sandbox flags, CSPEE
...: recently decided to have URN generating APIs emit a config aliased with some unique identifier
...: some properties are opaque to embedder, some not (e.g. width, height)
...: instead of a URN, getting a config object that might get a WebIDL interface
...: plug it into the fenced frame config IDL attribute
...: in this world, no src attribute, but would have a FencedFrameConfig object instead
...: lets you read those transparent properties
...: for non-opaque URL cases, can construct with an ordinary URL
...: more ergonomic than an opaque URN plus APIs to map it to data
...: simpler lifetime, don't need to worry about postMessage, etc

Dom: URL restrictions
...: was originally potentially trustworthy URLs (and URNs that map to that)
...: now confined to https:, about:blank, localhost; didn't want to deal with data: and blob: URLs associated with potentially trustworthy origins

Dom: CSP
...: new fenced-frame-src directive, applies normally to non-opaque URLs
...: for opaque case, CSP reporting would be a data leak
...: can only specify scheme source and host source, with wildcard in hostname and port

Dom: user activation, focus
...: don't propagate across fenced boundary in either direction
...: for user activation, clicking a fenced frame does not user-activate the embedder (unlike cross-origin iframes, where this propagates a transient user activation)
...: same for a nested fenced frame
...: same-origin documents across the fence also don't share user activation (e.g. A-B-A ancestor chain)
...: programmatically setting focus across the boundary is a communication channel
...: script inside fenced frame can only request focus with transient user activation
Christian: if it doesn't consume user activation, can't you use that to communicate?
Dom: [specific scenario too detailed to scribe]

Dom: session history
...: not part of joint session history
...: only replacement navigation for the entire fenced frame subtree
...: also fully isolated, can't observe history.length outside the fence (it's always 1 inside the fenced frame tree)

Dom: Spec
...: spec is mostly a shell right now
...: https://bit.ly/browsing-contexts
...: nested browsing contexts participate in frame tree, can see window.top, window.parent, etc., accessible by name, inherits certain properties
...: whatwg/html#6315 session history rewrite (earlier breakout session) provides primitives we intend to use to spec fenced frames
...: PR describes a "top-level traversable" which manages a session history and spans browsing context swaps
...: fenced frames will be the first nested traversable which has full frame tree isolation and doesn't participate in embedder's browsing context group

Dom: browsing context groups and SPECTRE
...: assumption: BCGs can always be in another process, even without site isolation (OOPIF)
...: fenced frames would be the first time we require two BCGs in the same process in such browsers
...: new invariant: all parent BCGs may be placed in separate processes, regardless of OOPIFs support
...: child BCGs (fenced frames) must opt into COEP, cannot be cross-origin isolated without parent

Alex Christensen, Apple: you mentioned it cannot do anything weird to session history; can it navigate?
Dom: it can navigate itself, replace-only
Alex: are there other parent browsing context groups you have in mind?
Dom: prerender, portals

Ben Savage, Meta: you talked about technical difficulties of creating this isolation; I'm sure you'll figure it out – how will end users understand that this part of the page can't communicate with the other part of the page? Perhaps one of the original motivations for

this is discomfort from creepy ads, not understanding why they're seeing the ad that they're seeing (due to interest group based targeting). I don't believe end users will conclude that an on-device auction opaque to the site occurred. Users will assume the website knows something about them. Is there some way we can visually prove this to users? We've missed the main problem.
Dom: Main problem is privacy more than creepiness. We could preserve privacy even while showing more specific information, even though that's a non-goal.
Michael: You're asking what the goal of Privacy Sandbox is. In order for PS to succeed, we want the web to be more private *and* feel more private. Users should actually have their privacy protected, and feel that way. Neither alone is sufficient. Here we are focused with fenced frame on actually being more private. There's also a UX problem to convince people that even though they're seeing targeted ads, that doesn't mean everybody learns stuff about them. You could imagine an ad campaign or similar. That's a huge problem, but it's not limited to technical things like what is displayed on the screen along with the ads. The scope also includes how the user perceives the guarantees the device, OS, etc promises.
Ray Cromwell, Google: For some vendors, users will tend to believe the vendor's marketing that this is true. But not necessarily for other vendors, unless they really put effort into it. It might work eventually, but certainly not right away. It has to be an educational process.
Michael: I don't think you'll be able to tell by looking at the ad that it is in a privacy-safe box. We have a bigger picture, the entire privacy stance of the web platform. We will need to communicate that we've improved it. Possibly not even limited to the web platform, but also across the entire device. It's a human problem, not a technology problem.
Ben: We've struggled to explain even basic things about technology, like the difference between Chrome and Safari. I find this argument unconvincing. I think your definition of privacy is too narrow – it means that the way data flows aligns with the way you expect data to flow. You seem to be saying that people who have seen the marketing materials will be convinced.
Dom: I don't think the UX concern is out of scope, but the bits over the wire carrying your information is the most urgent problem.
Ray: Ethically, you should make it secure and privacy, even if the user doesn't know or care. Eventually we decided everything should be HTTPS by default.
Alex Cone: [missed one point, something about this being only one part of the solution] Meta could get a lot of benefit from Google, Apple, etc doing this because you're going to need to confront this problem, too.
Charlie Harrison, Google: Fenced frame proposal has legitimate user-facing privacy wins in the form of restricting what inputs can go into it. We shouldn't dismiss that. K-anonymity check, etc, will change the kind of content put in front of users.

Alex Cone: Also won't see 200 different network requests when ads load.
Charlie: having a new element for displaying this kind of content creates an opportunity for more user control
Brian May: you're suggesting in the future that users will know that using a particular browser guarantees privacy in ads
Michael Kleber: the goal of Privacy Sandbox and other efforts discussed here is not a product-specific thing about a particular browser, but making the web platform inherently more private across browsers and platforms. As the people who build the web, we will want to communicate to the web users that the web is substantially more private than it used to be.
Brian: Today people assume that the author of the page owns the page and its contents. The fenced frame makes it harder for the user to know who to complain to.
Michael: This is already true today.
Brian: Today publishers know what ads are shown on their site. This would make the publisher not know what's shown on their site.
Michael: That's not true. Publishers can exert all the same types of control that they can today over what things might appear, and allows publishers to know after the fact, in aggregate, what sorts of things did appear. In the long term, once we remove a temporary carveout [didn't catch detail], publishers would not be able to build up a profile of all the ads that a particular user saw while visiting the site. That kind of profile building wouldn't be possible with new information flow. But knowing what ads appeared on my site – they'd have the same capabilities.
?: Is this supposed to be transparent to the user? Would they actually see different ads? Would this remove retargeting ads from the web, or the same ads, just more privately?
Michael Kleber: This is more of a question about FLEDGE than fenced frames. If you want to talk about the FLEDGE API, come to the WICG session on Friday. FLEDGE does impose some restrictions on the ways in which ads can be targeted. Users may notice a difference in the ways in which ads are chosen. These are two separate problems – we need to stop people from doing the tracking which is possible today (fenced frames partially solves this), and we need to make the future of ads feel less creepy. Restricting the flow of information is necessary but not sufficient to do that. It's not only browsers and platforms that need to do that. If ad tech wants to creep users out, they can do so. We can make it harder, but we cannot completely solve it, just push it in a good direction.
Simeon Vincent, Google: Getting back to the UX question. This revolves around adoption. What motivates authors to use fenced frames over iframes or raw HTML?
Dom: To embed ads based on cross-site identity, it won't be possible without fenced frames.

Stephen McGruer, Google: also necessary for certain payments use cases
Simeon: So the motivation is that they make more money from their ads?
Dom: For ads specifically, yes.
Simeon: Part of the challenge is having websites convince people that they are properly handling user data, not the browser itself.
Ben Savage, Meta: Yes, fenced frames and FLEDGE are separate, but the incentives are entangled. With regard to fenced frames specifically, I'd like to know which browsers have stated opinions and are there multiple implementers for both FF and FLEDGE?
Dom: For FF, we've talked to Apple (but didn't request a formal position), with a positive discussion on a previous version. Will request formal positions soon. Have gotten some positive feedback (but not formal position) from Mozilla. In Chromium we've had two implementations, one based on iframe and shadow DOM, one with a more radical architecture change. We think the fact that this is possible suggests other vendors will be able to implement.
Shivani: Microsoft Edge has a PARAKEET proposal which also uses fenced frames.
Michael: wrt FLEDGE, similar answer. Ongoing bilateral efforts to unify FLEDGE and PARAKEET over the last year. That's what the Friday session is about. Apple seem interested in this as a way of displaying things, but have concerns about the way of choosing ads in FLEDGE, and in particular the fact that it involves off-device stuff like trusted servers. We don't know if they will find some alternative as a "happy medium". Mozilla has been focused on other things and has no expressed position.
Shivani: we focused on other use cases with Apple, and we don't know their stance on the latest proposal