October 16, 2015

Federal Trade Commission
600 Pennsylvania Ave NW
Washington, DC 20530

Re: Cross Device Tracking Workshop

To whom it may concern,

The Future of Privacy Forum (FPF) is a think tank seeking to advance responsible data practices and is supported by leaders in business, academia and consumer advocacy.[i] We thank the FTC for providing this opportunity to comment on cross-device tracking and submit the attached report in response to the FTC's workshop on November 16, 2015.

Sincerely,
Jules Polonetsky, Executive Director
Future of Privacy Forum

[i] The views herein do not necessarily reflect those of the Advisory Board or supporters of the Future of Privacy Forum.

Cross Device: Understanding the State of State Management

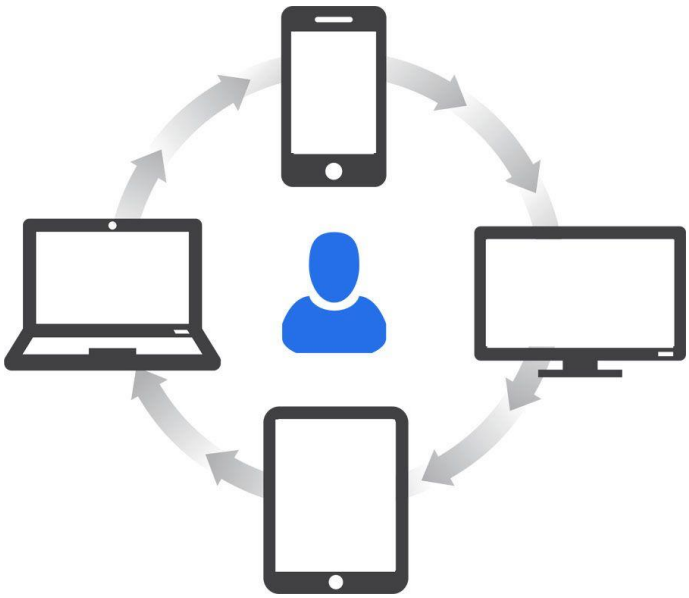

FUTURE OF PRIVACY FORUM

Jules Polonetsky, Executive Director

Stacey Gray, Legal & Policy Fellow

October 16, 2015

Contents

I. Why They Track[1]

For nearly 20 years, ad networks have used cookies to track internet users to tailor the ads they see across the Web. But despite the ubiquity of this kind of behavioral advertising, the practice continues to be a source of consumer policy debate. Advertising trade groups claim that the practice provides users with ads that are more relevant and useful. Critics contend that tracking uses private information in an attempt to unfairly manipulate consumers, and express concern that the creation of profiles could be used to discriminate against individuals who visit sensitive sites. Some leading web publishers worry that they don't have enough control over the practices of ad tracking that third parties routinely conduct on their sites, or that behavioral targeting commodifies their previously premium audience.[2] Other publishers are grateful for the support that ad revenues provide for their services and content. Efforts such as Do Not Track,[3] which attempts to provide new consumer privacy controls for tracking and ad targeting, have failed to reach consensus largely because many of the most ardent stakeholders hold diametrically opposed views about whether targeted ads are an essential good or an undesirable practice.[4]

Hidden from the acrimonious debate over behavioral advertising are the less provocative reasons that websites work with tracking companies. By placing cookies on users' browsers, publishers can better understand how many unique users visit a website or see an ad delivered across many different sites. They may "cap" an ad to make sure that each user sees a giant pop-up ad only a certain number of times. And advertisers can learn which ads are most effective at bringing users to their website.

It is now fairly commonplace that advertisers track users to analyze and understand consumer behavior and trends. Public scrutiny has encouraged business to be proactive when it comes to respecting data from consumers. Some companies go to great lengths to anonymize the data they collect with encryption and double-blind processes through third-parties to protect the privacy of consumers. Rarely do these privacy principles detract from a great ad campaign: leading companies regularly publicize case studies to underscore their advertising prowess.

Nevertheless, critics of behavioral tracking persist. Many point out that TV, radio and magazine advertisers use research panels of volunteers to measure ad effectiveness. Why do websites need to measure effectiveness with greater precision, given the complaints about tracking? But these other media can rely on the power of their message which is supported by sound, pictures, and emotional stories. Users can recall the good ads, and the best become part of pop culture. "Good to the last drop" -- Maxwell House. "If I were an Oscar Meyer Weiner." "Mikey likes it!" The tiny Web banner ad can hardly compete with TV, radio and magazines except for in one way: in being precisely measurable.

Furthermore, the panels used to measure TV and radio rely on the ability to sign up a small pool of willing representative participants who have their activity tracked and extrapolated to represent the

---

[1] This section adapted from Jules Polonetsky & Christopher Wolf, *Why They Track Us*, HUFFINGTON POST (May 4, 2012), http://www.huffingtonpost.com/jules-polonetsky/why-they-track-us_b_1475027.html.
[2] *See, e.g.*, Jason Kint, *Behavioral Advertising Might Not Be As Crucial As You Think*, ADVERTISING AGE (Feb. 25, 2014), http://adage.com/article/datadriven-marketing/behavioral-advertising-crucial/291858/.
[3] Do Not Track: Universal Web Tracking Opt Out, http://donottrack.us.
[4] *See generally*, Omer Tene & Jules Polonetsky, *To Track or 'Do Not Track': Advancing Transparency and Individual Control in Online Behavioral Advertising*, 13 MINN. J. L. SCI. & TECH 1 (2012).

large media audiences. Some online panels are available for the large web properties, but would leave the many hundreds of thousands of small web publishers with no way to represent the make-up of their audiences or represent the performance of their ads. And so, although industry views about the value and merits of targeted ads are diverse, wide business consensus considers the cookies and web beacons of web tracking and measurement to be essential.

In these comments, we describe how and why web cookies are becoming less effective as a way to enable this sort of online tracking, and discuss the alternative methods currently in use.

II. The Rise and Fall of Cookies

After their introduction in the 1990's, cookies quickly became the standard method for websites to engage in "state management"—the ability to remember information about the same user over time. In the early years of the Internet, web sites contained mostly static information, like pages in a book, and were unable to save user preferences or display personalized content.

The solution was identified by Lou Montulli of Netscape,[5] who realized that if a website stored a small piece of data—a "cookie," typically a short string of unique text—onto the user's hard drive, then the next time the user visited that site (via the browser making an HTTP request to access it), the site would recognize that cookie on the user's device. This allowed websites to distinguish between new visitors and returning visitors (thus permitting more accurate visitor counts), as well as to identify an individual as the same person who previously visited certain sections of the site, clicked on specific advertisements, or added specific items to a shopping cart.

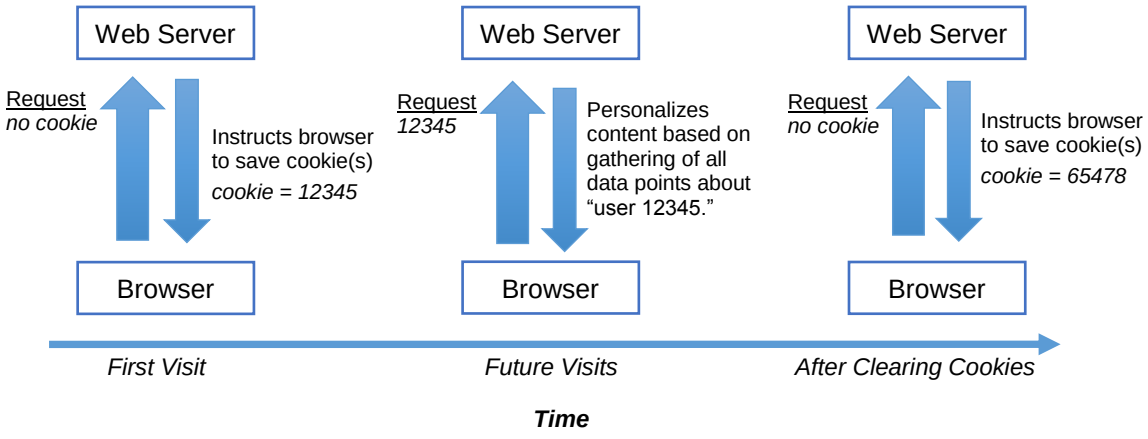

Cookies have revolutionized the Web for both consumers and publishers by allowing personalized interaction with websites over time—but they only work effectively for as long as the individual chooses to retain them and continues to use the same device and browser. For an average website that seeks to adapt to a user's preferences without requiring her to log in, the cookie is only useful

[5] *See* U.S. Patent No. 5,774,670, *Persistent client state in a hypertext transfer protocol based client-server system* (filed Oct 6, 1995).

as long as the user is primarily and consistently using the same device (e.g. their home desktop) and the same browser to engage with that website.

Today, the proliferation of connected devices is causing the traditional cookie model to become less effective. The modern user now interacts with media using a growing list of separate internet-connected devices throughout the day, including not only her smartphone, laptop, and tablet, but now also TV, watch, fitness tracker, and connected car. Moreover, consumers access the web via multiple Internet Service Providers (ISPs): home, office, mobile, fixed line, and Wi-Fi. Because these devices and systems are not directly connected, this growing use of smart devices has created a real challenge for advertising and marketing industries.

A. Exponential Increase in Third Party Cookies

When the initial design for the cookie was introduced in 1995, cookies were envisioned as pieces of data that would be primarily delivered from a website's own domain to the user's hard drive for the purpose of personalizing that particular domain, i.e. "first party cookies."[6] The most obvious example at the time was the shopping cart, permitting a website to remember what an individual had placed into an online checkout and prompt her to complete payment when she returned later. Similarly, a weather site might save a user's inputted zip code in a cookie in order to consistently display the weather for that geographical area, or a commercial website might target their discounts to an individual who visited certain pages. These types of first party cookies were straightforward in purpose, and the number was generally minimal.

Very soon, however, the use of cookies became more complex. Web publishers began to dispatch cookies from large numbers of third party systems, including advertisers, content providers, and marketing trackers. This growth of third party involvement, as well as the added complexity of the cookie mechanisms, can leave a user's browser with hundreds of cookies that are used to track, analyze, manage, and target content, advertising, and other features.

This increase in third-party cookies affects different industry players in different ways. For web publishers, there is the concern that an increasing reliance on mechanisms that deploy third party cookies slows down website load times for their users. Some publishers also cite concerns about data leakage—when a third party collects data about users and then uses that data without the original publisher's permission—as a reason to restrict those mechanisms.[7] However, web publishers may also risk losing advertising revenue, and thus the ability to provide content to users at free and lowered costs, if they do not support third party mechanisms.

For consumer privacy advocates, the exponential increase in the use of cookies has caused concern that third party companies have made cookie-linked data available via data exchanges, resulting in widespread data dissemination and diminished user control over their information. Yet, if concerned users decide to disallow or delete cookies, this further diminishes their usefulness as a

---

[6] *See generally*, John Schwartz, *Giving Web a Memory Cost Its Users Privacy*, NEW YORK TIMES (Sept. 4, 2001), *available at* http://www.nytimes.com/2001/09/04/business/giving-web-a-memory-cost-its-users-privacy.html.

[7] *See, e.g.*, Garett Sloane, *Google Cracks Down on Marketers' Access to Data*, ADWEEK (Oct. 3, 2014), *available at* http://www.adweek.com/news/technology/google-cracks-down-marketers-access-data-160543.

way to provide consistent personalized content across devices and logins, including advertising as well as consistent privacy preferences.

B. The Democratization of Data

In early days of web advertising, a key goal of ad networks was the aggregation of as much data as could be used to make the most effective ad targeting decisions. Leading ad networks boasted about the breadth of their networks, the partners who shared data with them, and the third party data that they had linked to cookies. In 1999-2001, the merger of DoubleClick, a leading advertising network, with Abacus, an offline data collector (sparking a wave of privacy concerns) was driven by DoubleClick's desire to expand its network by adding the rich data from the Abacus data coop to its own data on web surfing patterns. For ad networks, assembling and linking data on a range of demographic, psychographic, and purchase history information was an expensive and technologically complex endeavor.

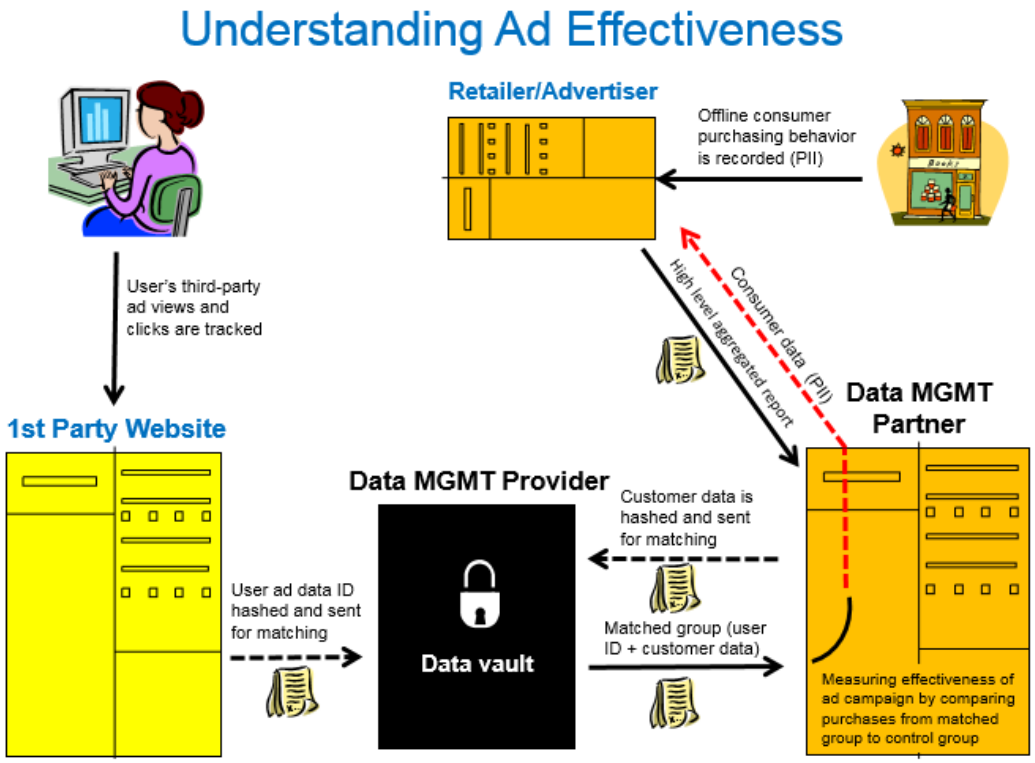

Today, in contrast, data has been "democratized." Advances in technology have lowered the costs of storing and managing data, and as a result, individual consumers and small businesses have direct access to unprecedented amounts of data about themselves and others. Companies large and small can efficiently link their own customer data and many dozens of massive data sets using self-service tools or by working with data exchanges. All of these data sets are linked to cookies and widely available for targeting and tracking.

Companies selling their data for online appending via the BlueKai data marketplace

C. Disconnect between Devices

Increasingly, users now access the internet from a diverse range of internet-connected devices. However, because cookies are specific to each unique device and browser, web publishers and third parties may not recognize that the same person is behind each device and browser. When a user visits a website from his laptop to check, for example, basketball game statistics, a cookie is delivered onto that laptop. However, if he uses his phone later to check the same content, a new cookie is set, as the publisher does not recognize him as the same visitor. This lack of connectedness constrains web publishers from delivering customized and consistent content, services, and features, as well as from enabling the tracking and customized advertisements that often allow the content to be provided to the user for free or at a reduced cost.

D. Divide between Mobile Web Browsers and Mobile Apps

In addition to the divide between devices, there is a lack of communication between mobile web browsers and between mobile apps that constrains tracking even on the same device. While mobile web publishers and content providers may use cookies (to the extent permitted by the mobile web browsers), mobile apps do not. Without a web browser intermediary between publisher and user, an app cannot place cookies into storage on the device; instead, it must rely on the device's manufacturer-provided identifier.

Initially, app developers and other third parties tracked user behavior in apps using a range of operating system identifiers, device identifiers, MAC addresses, and other identifiers assigned by manufacturers or operation systems and permanently linked to the device. This generated privacy concerns from advocates who criticized the use of identifiers that were fixed and that could not be controlled by users.

In response to concerns over the use of device identifiers by third parties, mobile platforms such as iOS and Android began replacing this practice with new advertising identification numbers (e.g. the Apple IDFA, and the Android Advertising ID), which could be re-set by the user.

Mobile Apps

- Apps do not support cookies.
- Initially, third parties tracked apps using device identifiers.
- Platforms replaced with Apple **IDFA** and Android **Advertising ID**
- No cookie syncing needed, mobile app IDs are global

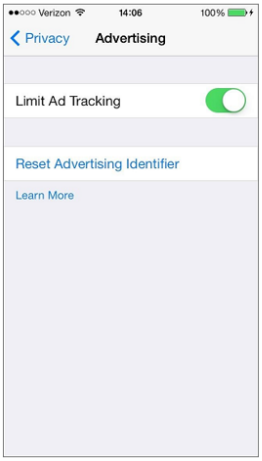

E. Safari

Although the Safari browser on Apple desktop computers blocks third party cookies, Apple's market share of the desktop market has always been relatively small. But on the web, Safari has a 39% market share as of September 2015 (down from 45% nearly a year ago due to the rising prominence of other web browsers).[8] As a result, Safari's practice of blocking third-party cookies on its mobile browser diminishes the ability of advertisers to track a user who primarily uses her mobile device to browse the web.

Because mobile Safari users are often considered to be more affluent than competing device users, Safari's mobile ad-blocking makes this attractive audience difficult to track and threatens ad tracking business models.

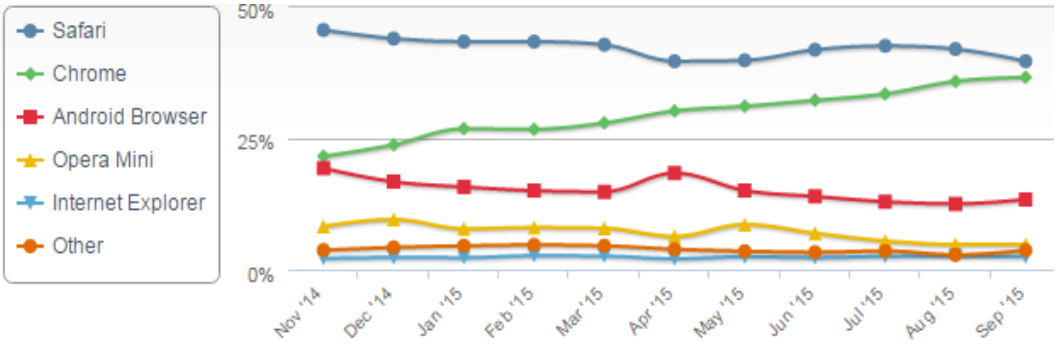

The recent addition of ad-blocking capacity in iOS 9 may further diminish the ability of web publishers and marketers to track users efficiently.

[8] NetMarketShare, Mobile/Tablet Top Browser Share Trend, https://www.netmarketshare.com/browser-market-share.aspx?qprid=1&qpcustomb=1 (last accessed Oct. 13, 2015).

III. Emerging Alternatives to Cookies

There are a number of emerging ways that are being employed to enable tracking of the same individual across devices and platforms.

A. Device or Client Identification (Deterministic Matching)

Deterministic matching systems recognize whether two users are the same based on at least one unique personalized data point, such as a name, email address, home address, or phone number. The most recognizable way that this occurs is when a user logs in to an authenticated site, such as Google or Facebook, on multiple devices—say, a mobile app, and a desktop—at which point, the company and its subsidiaries can recognize these authentications as arising from the same person.

Thus, for companies with users that "log in" regularly on different devices, maintaining state with their users across the ecosystem is relatively simple. This also means the data linked to web cookies set by the authenticated site can be linked to the user's corresponding mobile identifiers. These companies can also assist their ad network divisions to more effectively target a unified user.

Another form of deterministic tracking across devices occurs in the world of third party ad networks. A third party advertiser can match the cookie of a user who has provided her email address to a partner website, with the mobile activities (via Mobile Ad ID) of a user who has provided that same email address to any partner app. These third party companies typically take steps to hash names and email addresses in order to avoid restrictions that partners may have around sharing personal information. Hashing, a process of translating personal data into corresponding (but shorter) strings of randomized characters, permits a data matching partner to receive data from different sources and match it without directly viewing the personal information underneath.

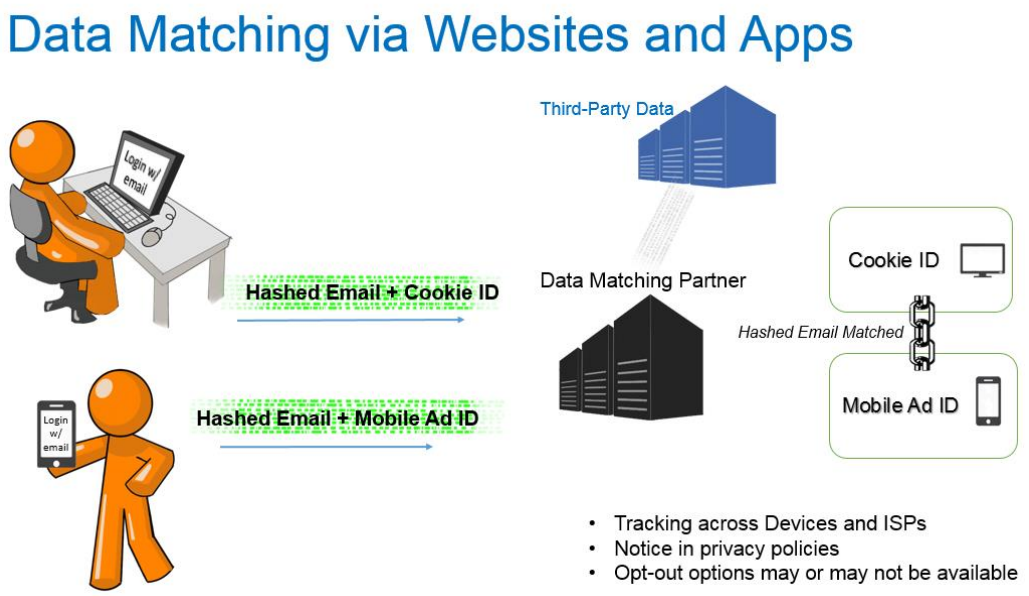

As visualized above, a data matching partner can receive hashed data from an array of difference sources in order to match users across devices. For example, a hashed email from Website A (and associated web cookies) can be matched by a data matching partner with the same hashed email from Mobile App B (and that device's associated advertising ID).

B. Statistical Identification of Devices (Probabilistic Matching)

In the absence of authenticated data, it is often still possible to track individuals, with different measures of accuracy, across devices and platforms on the basis of statistical information gathered from the device, browser, app, and operating system. For example, this information can include the fact that multiple devices (say, a laptop and a phone) consistently use the same home Wi-Fi router, and are turned on at roughly the same time every evening. Using these kinds of rough data points, and many others that may be collected about a device, a system can infer—within ranges of confidence—that those devices are being used by the same person.

This kind of systematic inference-making (probabilistic matching) can create one matched identifier that links apps and programs both within and across devices. Unlike cookies, however, statistical identification is based wholly on probability, and identification algorithms that promise greater accuracy (say, 90% confidence that two users are the same person) are more valuable to marketers and advertisers than algorithms with lower accuracy.

One common example of probabilistic matching occurs when an individual uses multiple devices on the same home Wi-Fi router. This enables the collection of the user's IP Address, which can be matched to cookies and mobile advertising identifiers to provide cross-device tracking as well as geographically targeted content.

Ad Networks using Home Wi-Fi

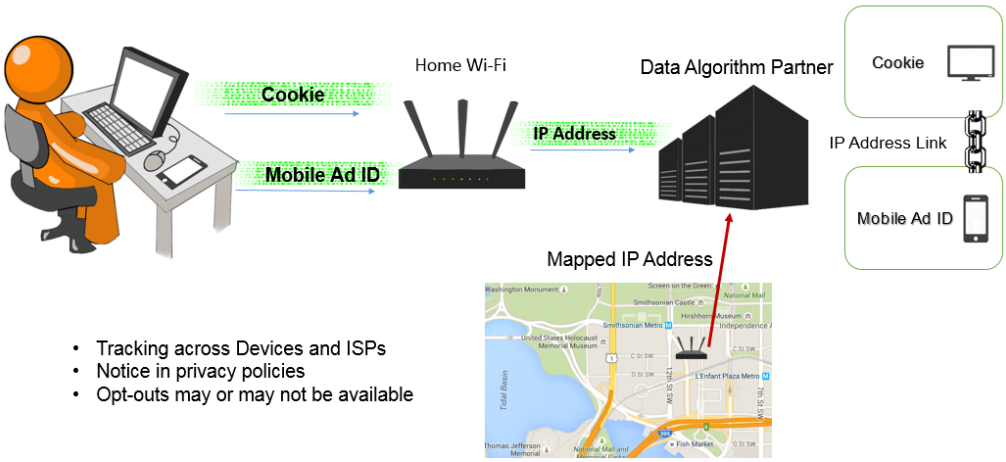

- Tracking across Devices and ISPs
- Notice in privacy policies
- Opt-outs may or may not be available

Another method of probabilistic matching takes advantage of the fact that browsers and devices, even when they don't have an available cookie or advertising ID, nevertheless have unique attributes that allow the creation of a statistical identifier or a browser "fingerprint". This method, often also known as server side recognition or device tracking, can typically identify the characteristics of a browser with great accuracy.

Browser Fingerprinting

• Browser is queried for its agent string, screen color depth, language, installed plug-ins with supported mime types, time zone offset and other capabilities, such as local storage and session storage.

• 18 bits of entropy, meaning that only one in 286,777 other browsers will share its fingerprint.

• Industry terminology: probabilistic targeting, server side tracking or device recognition.

• For mobile, time differential latency also used.

navigator.userAgent *// "Mozilla/5.0 (X11; Linux i686) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/27.0.1453.110 Safari/537.36"*

navigator.language *// "en-US"*

var **plugins** = $.map(navigator.plugins, function(p){ var mimeTypes = $.map(p, function(mimeType){ return [mimeType.type, mimeType.suffixes].join('~'); }).join(','); return [p.name, p.description, mimeTypes].join('::'); }); $.each(plugins, function(i, p){ *// truncate only for blog example* if

(p.length > 80){ console.log(p.substring(0, 77) + '...'); } else{ console.log(p); } }); */* Shockwave Flash:Shockwave Flash 11.7 r700:application/x-shockwave-flash~swf,a...*

Chrome Remote Desktop Viewer:... Widevine Content Decryption Module:Enables Widevine

*licenses for playback of ... Native Client::application/x-nacl~nexe Chrome PDF Viewer::application/pdf~pdf,application/x-google-chrome-print-prev... Google Talk Plugin Video Accelerator:Google Talk Plugin Video Accelerator ver... Google Talk Plugin:Version: 4.0.1.0:application/googletalk~googletalk Google Talk Plugin Video Renderer:Version: 4.0.1.0:application/o1d~o1d Shockwave Flash:Shockwave Flash 11.2 r202:application/x-shockwave-flash~swf,a... */* screen.colorDepth *// 24* new Date().getTimezoneOffset(); *// -240* !!window.**localStorage** *// true* !!

window.**sessionStorage** *// true*

C. Internet Service Provider State Management

In addition to these methods, Internet Service Providers (ISPs) can enable tracking of users across devices by inserting a unique identifier in web traffic that can be used by an ad network partner as a cookie alternative. This method is not feasible when web traffic is encrypted, or when the user is on Wi-Fi or using another ISP (for example, at home vs. at the office).

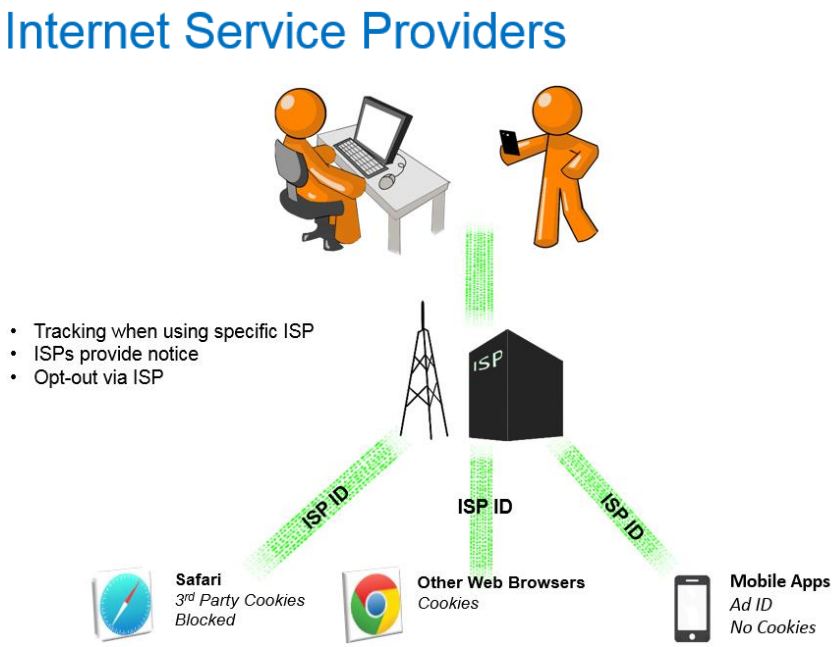

D. Probabilistic + Deterministic

Given the fact that deterministic tracking provides greater certainty but limited reach, and probabilistic tracking provides broad reach but less reliable certainty, it should be no surprise that many companies are now using a mix of both methods.[9]

E. Emerging Co-Op Model

Increasingly, some publishers and content providers are choosing to enter into "co-ops," or lateral agreements with each other to share authenticated personal data. In other words, rather than providing data to third party advertisers, partner companies can directly exchange their databases of linkages between users and devices and browsers or work with a common partner who serves

[9] *See* Allison Schiff, *Oracle Partners With Tapad – Because Probabilistic Vs. Deterministic Data Isn't An And/Or Sort Of Thing*, ADEXCHANGER (Oct. 15, 2015), http://adexchanger.com/data-exchanges/oracle-partner-with-tapad-because-probabilistic-vs-deterministic-data-isnt-an-andor-sort-of-thing/.

as a trusted intermediary to enable each participant to more broadly recognize users.[10] Typically, companies share this data in hashed format, in order to minimize privacy concerns.

F. Location Targeting Across Devices

One key goal of marketers is to understand whether a user who was exposed to an *online* advertisement later traveled to a physical retail location. Since mobile operating systems and apps routinely request permission from users to obtain location information, those companies can link location to mobile identifiers—and then via cross device methods to a cookie. Now, an ad network with access to this data can report that the ads it displayed on a desktop computer were indeed responsible for actually bringing users into a store.

Another method allowing apps to collect location information is the use of "beacons" —small devices consisting of a chip and other electronic components (e.g. antenna) on a small circuit board. Beacons are essentially radio transmitters that broadcast one-way signals to devices that are equipped to receive them. These devices allow the mobile app to determine (typically via BlueTooth) a user's location in proximity to the beacons, which may be installed at various places throughout a retail location, such as in front of a special display of products.[11]

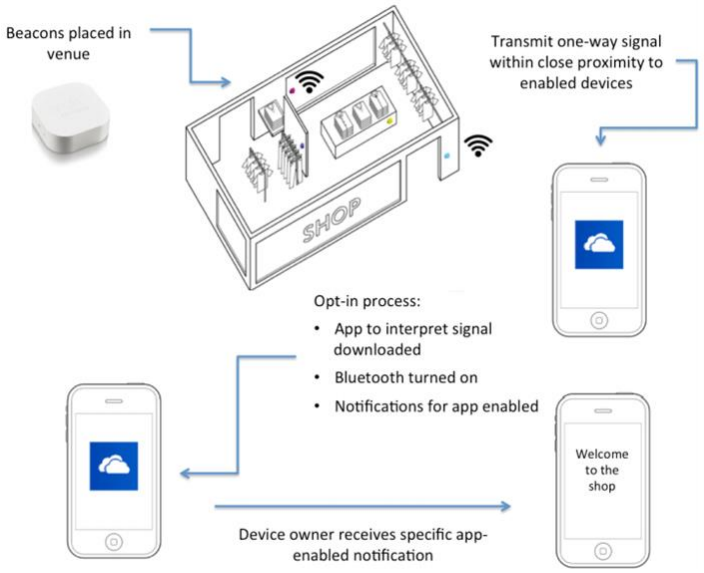

[10] Adobe recently announced its interest in forming data co-ops. *See* Press Release, *Adobe and Publicis Groupe Team Up to Deliver First Always-On Global Marketing Platform* (Sept. 10, 2014), *available at* http://www.adobe.com/news-room/pressreleases/201409/091014AdobeTeamsUpWithPublicisGroupe.html.

[11] *See generally*, GREG STERLING, JULES POLONETSKY & STEPHANY FAN, FUTURE OF PRIVACY FORUM, UNDERSTANDING BEACONS: A GUIDE TO BEACON TECHNOLOGIES (Dec. 2014), *available at* http://www.futureofprivacy.org/wp-content/uploads/Guide_To_Beacons_Final.pdf.

G. Benefits and Risks

Supporters of cross device tracking argue that it is necessary to maintain the tracking and targeting that supports their business models and web content. Smaller businesses argue that cross device tracking methodologies are needed to enable them to compete with larger consumer facing brands that have hundreds of millions of authenticated users who can be recognized across devices, apps, and browsers. Detractors of ad tech see it as another unfair extension of a business model of which they disapprove.

Historically, cross-device tracking emerged, in part, as a viable solution to serious problems of cybersecurity and fraud detection. While these comments are focused on the benefits and risks of cross-device tracking within the advertising technology ecosystem, it should be noted that the use of cross-device tracking for user security, anti-fraud, anti-spam, authentication or similar uses should be non-controversial.

i. Transparency

In the early days of tracking and targeting, transparency was a simpler affair. Web sites could disclose targeting in privacy policies, and link to the policies of ad networks or central industry disclosure and opt-out sites. As the proliferation of third party trackers became more intertwined and complex, sites lost visibility into what third parties might be in the chain of companies passing data along until an ad is selected and delivered. Publishers and advertisers have trouble representing the details of tracking technologies used by intermediaries when they often don't know who those intermediaries are. The hyper efficiency of the ad market has created transparency challenges that make it difficult to accurately describe or commit to specific technologies or policies.

Furthermore, when the advertising market was reliant on cookies, users could be reliably advised that deleting or blocking cookies provided a strong set of controls over tracking. As methods that do not rely on cookies are adopted, a gap is created between the privacy tools provided by browsers or operating systems and the realistic ability of consumers to control new tracking methodologies.

To some extent, central industry disclosure and opt-out sites provide a response to this issue, as parties that license the standard tracking icon can point to a disclosure and a location that provides users with options to decline to be tracked by participating companies. But transparency provided by device settings or browser settings is increasingly limited. ISPs can offer disclosures to their customers, but today many users rely on Wi-Fi, and use multiple ISPs at home, work, and on devices, limiting the ISP's reach.

ii. Effectiveness of Scope and Controls

The Do Not Track (DNT) debate illustrates the gap between privacy advocates and staunch ad targeting champions. Ardent privacy advocates argue that consumer controls over tracking should be robust and permit restriction even of web sites' ability to track over time other than for very limited security or functional purposes.

Industry actors argue that consumer choice should be limited to an opt-out of web surfing related targeting, leaving appended data targeting and tracking for the sake of measurement and reporting fully in effect. A spectrum of views exist in between these poles, resulting in policy paralysis. Very

few companies treat the Do Not Track signal even as a behavioral advertising opt-out.[12] An even fewer treat the DNT signal as championed by DNT advocates.[13]

Thus, it should be no surprise that the scope of any ability to opt out of cross device tracking has yet to be defined. Some companies can at least technically offer a firm ability to break the link between users' multiple devices, while for others the challenge is less technically feasible. Some do not believe it will be appropriate to break the link, but rather simply to turn off targeted advertising, while others support a broader opt-out which would de-link devices. Some believe users who opt out intend to do so across devices, other maintain that users express preferences solely for the device at hand.

IV. Conclusion

The online advertising world has moved from a competition to acquire data to a struggle to maintain state management. As technologies that provide unique identifiers expand to include wearables, home thermostats, beacons, smart lighting, and every type of device in the Internet of Things, tracking will increasingly include data from a broader range of sensors and collect even more intimate information about users—or simply individuals within range of a device.

As facial recognition and drones enter the debate, even broader data and tracking may be available. Unified tracking will continue to be challenged, and fully effective privacy preferences will similarly be difficult to achieve. If leading companies that authenticate users and provide platforms across many consumer technologies dominate, tracking will be more centralized and stable, and so may be the ability to provide controls. If users interact with many diverse parties, state management will be more fragmented, limiting wider tracking but also limiting the deployment of wider scope controls.

The complexity of the ecosystem has already made transparency a huge challenge. Privacy policies must either be high level and generic or technical and detailed, each difficult for the typical consumer who has long abandoned interest in scrutinizing these policies. Small screens and vanishing screens will only exacerbate this problem. Transparency may provide an important basis for FTC enforcement, but users may be increasingly challenged to fully grapple with these issues.

An increasingly central part of the privacy discussion is focusing on issues such as fairness, equity, power imbalances and discrimination. Data can be used to discriminate, intentionally or in hidden but injurious ways. In the words of FTC Chairwoman Edith Ramirez in 2014:

> "Big data analytics raises the possibility that facially neutral algorithms may be used to discriminate against low-income and economically vulnerable consumers . . . There is the worry that analytic tools will be used to exacerbate existing socioeconomic disparities, by segmenting consumers with regard to the

[12] *See* FUTURE OF PRIVACY FORUM, *All About Do Not Track (DNT),* allaboutdnt.org (last visited Oct. 13, 2015).
[13] *See* Press Release, ELECTRONIC FRONTIER FOUNDATION, *Stop Sneaky Online Tracking with EFF's Privacy Badger* (July 21, 2014), https://www.eff.org/press/releases/stop-sneaky-online-tracking-effs-privacy-badger.

> customer service they receive, the prices they are charged and the types of products that are marketed to them."[14]

But data can also be used to empower individuals and organizations.[15] Going forward, the solutions that seek to advance consumer protection and fairness will need to address the technologies used for tracking and assess their transparency, their controls and scope of controls, as well as the effects of the deployments of these technologies on different classes of users.

[14] *Big Data Discrimination: Is The Industry Responsible?*, ADEXCHANGER (Oct. 15, 2015), http://adexchanger.com/data-driven-thinking/big-data-discrimination-is-the-industry-responsible/.
[15] *See* FUTURE OF PRIVACY FORUM & ANTI-DEFAMATION LEAGUE, BIG DATA: A TOOL FOR FIGHTING DISCRIMINATION AND EMPOWERING GROUPS (2014), *available at* http://www.futureofprivacy.org/wp-content/uploads/Big-Data-A-Tool-for-Fighting-Discrimination-and-Empowering-Groups-FINAL.pdf.