Privacy and Tracking in a Post-Cookie World

A whitepaper defining stakeholder guiding principles and evaluating approaches for alternative models of state management, data transparency and privacy controls for consumers, publishers, and trusted third parties.

JANUARY 2014

© 2014 Interactive Advertising Bureau

This document has been developed by the Future of the Cookie Working Group in collaboration with the IAB's Mobile Marketing Center of Excellence.

About the IAB's Future of the Cookie Working Group: The Future of the Cookie working group is reimagining the technology used to identify consumers across multiple sessions and devices in a way that promotes greater persistence of both identity and user choice. The starting point is "Imagining a world where HTTP cookies were never invented." More information can be found at: www.iab.net/Future_of_the_Cookie_Working_Group

About the IAB's Mobile Marketing Center of Excellence: The IAB Mobile Marketing Center of Excellence, an independently funded and staffed unit inside the IAB, is charged with driving the growth of the mobile marketing, advertising and media marketplaces. The Mobile Center devotes resources to market and consumer research, mobile advertising case studies, executive training and education, supply chain standardization, creative showcases and best practice identification in the burgeoning field of mobile media and marketing. Our agenda focuses on building profitable revenue growth for companies engaged in mobile marketing, communications and advertising, and helping publishers, marketers and agency professionals understand and leverage interactive tools and technologies in order to reach and influence the consumer. More information can be found at: www.iab.net/mobile

IAB Contact Information:

Brendan Riordan-Butterworth
Director, Technical Standards
(212) 609-3734
Brendan@iab.net

Belinda J. Smith
Senior Manager, Mobile Marketing Center of Excellence
(212) 380-4720
Belinda@iab.net
mobile@iab.net

Executive Summary

This paper is provided to acknowledge and address the limitations of the traditional cookie for providing persistent user privacy choices and tracking in our evolving multi-device, multi-environment digital landscape (discussed in this paper the "state management" challenge). Presented here is an initial examination of the various solution classes of state-management technologies currently and potentially available, as well as their efficacy as measured by the needs of disparate stakeholders. This paper lays the foundations upon which best practices for implementation for each state-management solution class may later be constructed.

The scope of this work is to define guiding principles for stakeholders, evaluate each state-management solution class against these principles, and to educate the reader on the current state management landscape. The intention is not to champion a specific solution class over another, nor to mandate which path the industry should pursue to address the current state management challenges. Rather, the guiding principles set forth in this paper will serve as a consistent measure of current and future state-management mechanisms and solution classes.

Guiding principles are presented for consumers, publishers, and third parties; defining each stakeholder's needs and requirements from state-management mechanisms. The needs and requirements of each stakeholder group are given, irrespective of the current existence of a technology or solution class which fully addresses all of these needs.

1 Contents

2 Overview

Originally designed for temporary data storage, the cookie has long-since evolved into a fundamental infrastructure component of the Internet. However, for a variety of reasons, cookies are no longer an acceptable mechanism for “state management” (i.e. providing the information necessary for content creators and third parties to deliver personalized information and services to end consumers and respect their preferences for privacy, information transparency, and control). For online publishers the proliferation of cookies has slowed page load times, increased ad discrepancy counts, and led to concerns of data leakage. It has also perpetuated a broken compensation model, whereby publishers risk revenue loss if they don’t support third party cookies, as well as from users who block or delete cookies, and a tilted playing field favors large consumer website brands. Publishers also experience operational and privacy policy burdens as various privacy initiatives, browser defaults, and regulatory measures gain traction.

For online consumers the proliferation of cookies has increased anxiety in regards to their online privacy. Data collection is fragmented over many websites, devices, browsers, apps, etc.—making it exceedingly difficult for consumers to understand who may be doing what with their data and to apply privacy controls centrally and consistently, while ensuring these choices persist over time.

For third parties the reliance on cookies has resulted in a battle between a rapidly degrading economic model, and the costly, persistent, and high-volume deployment of cookies. Though cookies are increasingly ineffective as a state management mechanism the industry continues to deploy them at an escalating pace causing: excessive network traffic and related costs, “internet bloat,” regulatory threats, and anxiety among consumers and publishers alike.

In light of these challenges and their likelihood to intensify over time, the Interactive Advertising Bureau (IAB) and its Mobile Marketing Center of Excellence formed the Future of the Cookie working group to consider alternatives to the cookie. This analysis was grounded in a consideration of the needs and desires of online consumers, publishers and the third parties they trust.

3 What is State Management?

Imagine you work in a building with a security desk on each floor. Think how frustrating it would be if every time you walked into the building or went to a different floor you had to provide your name, company, job title, and ID so security personnel could make sure you're allowed to proceed. You would have to provide all of this information every time you left the building or went to another floor—even if you just went for a quick coffee break, or walked a guest to the elevator. To circumvent headaches such as these, security badges were invented. Now every time you enter your building or change floors you are able to swipe your badge at the security desk and that swipe provides information to quickly remind the system of all of your details and automatically gives you permission to proceed. Additionally, your security badge contains information about you that can only be read by the security desks in your building, it would not work if you swiped it anywhere else.

Similarly, "state management" refers to the method of and ability for systems to remember information about users, devices, or software applications over time. Cookies are the primary mechanism for state management on the Internet. They act like security badges for websites. Websites use cookies to remember things about the visitors they serve so that they can provide visitors more personalized content and services, and remember their preferences for future visits.

4 The Importance of State Management

State management enables publishers and third parties to deliver personalized content, advertising, and services to end users. This includes the ability to persistently adhere to end-user preferences for privacy, information transparency, and control. Publishers and third parties cannot perform these functions without state management. In today's connected world, with flourishing growth in digital content availability and choice, visitors have come to expect and value personalized digital experiences. Publishers, advertisers, and other third parties recognize personalization is critical to attracting, engaging and retaining desirable audiences, as well as honoring their privacy preferences. Many digital content and service providers also rely on revenues from personalized advertising to subsidize their cost of doing business. These subsidies can then be passed to the consumer in the form of free or substantially discounted content and services. An inability to capitalize on personalized advertising revenues, and adhere to consumer privacy preferences, could impact this subsidy chain making digital content and services more expensive for consumers. That is, consumers' current ability to access a myriad of digital content, information, and services—and to access it at little to no cost—may be tangibly reduced.

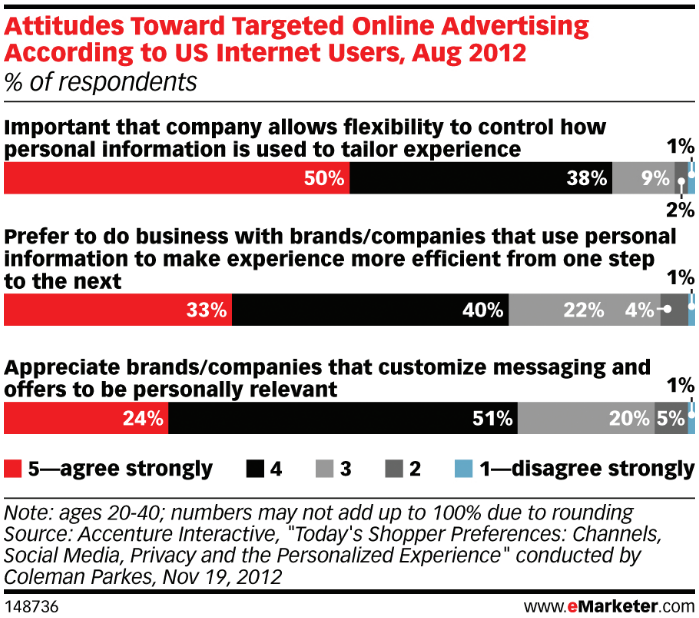

Consumers want more personalized content, but also want flexibility and control when deciding how personal information is used.

5 Why the Cookie May Be Crumbling

The current cookie approach to state management is fundamentally at risk for two main reasons:

1. The proliferation of cookies along with the resulting technical and privacy challenges

2. The growth and increased diversity of Internet-connected devices.

5.1 The Proliferation of Cookies

The cookie is the state management mechanism most commonly used today to support the many functions of digital personalization, reporting, and advertising. It is not owned or licensed by one party—rather, it is part of Internet Standards, and as such the Internet industry as a whole has innovated on top of it.

When cookies were originally conceived, nearly all of the content served to a web page was delivered from the website's own domain. For this reason, the number of cookies deployed was minimal and most were first party cookies. Websites today have become much more complex and sometimes use hundreds of third party vendors and systems in concert to deliver personalized content, services, advertising, features, and functionality.

Given the domain level access at the core of cookie functionality, each of these third parties typically deploys at least one unique third party cookie per domain, meaning hundreds of cookies could be stored on a visitor's browser after visiting a few web pages. Since the cookie is the primary foundation for so much of the Internet's functionality—managing user preferences, analytics, shopping carts, content recommendations, and advertising—cookies are being deployed by more companies, for more purposes, which is contributing to an exponential increase in cookies stored on visitor's browsers.

Online publishers find themselves stuck between a rock and a hard place. The proliferation of web beacons and pixels (mechanisms used to deploy third party cookies) slow down their website load times for visitors, increase advertising discrepancy counts, and lead to concerns of "data leakage". However, to remain competitive and attractive to visitors and advertisers, publishers have increasingly relied on multiple third parties to provide visitors with enhanced functionality and features. This gives publishers the option to risk revenue loss if they don't support third party pixels; though the proliferation of third party pixels has caused visitors to increasingly block or delete cookies (known as "cookie churn"), which also causes publisher revenue loss.

For online consumers the proliferation of cookies has increased anxiety over online privacy, transparency and control. With so many cookies being deployed, by sometimes unknown third parties, consumers are increasingly concerned about what tracking is occurring and by whom. As users become more aware of the data trail created as they surf across the Internet, but lack a fundamental understanding of how that information is used, many users choose to opt-out of tracking altogether.

With consumer concerns comes the very real prospect of regulatory intervention. Regulators are taking a close look at current practices and considering legislation to address consumer demand for increased transparency and choice, such as the FTCs recommendation for a "Do Not Track" mechanism[1]. As the appetite for intervention grows, the digital advertising industry faces increasing operational and compliance costs as regulatory measures become reality. To avoid these burdens on all sides, the industry is searching for solutions that can ease consumer and regulator concerns while proactively addressing current state management needs. Major browsers have also used "Do Not Track" settings (Internet Explorer) or are considering blocking all third party cookies (Firefox) as a mechanism for showing alignment with consumer's concerns.

[1] Federal Trade Commission. (2010, December 01). *FTC Staff Issues Privacy Report, Offers Framework for Consumers, Businesses, and Policymakers. Retrieved from ftc.gov: http://www.ftc.gov/opa/2010/12/privacyreport.shtm*

5.2 Cookies and the Diversity of Internet-Connected Devices

In the past most consumers accessed the Internet exclusively through a browser on a desktop computer. Today, however, consumers are accessing the Internet from a broad array of Internet-connected devices including phones, tablets, gaming consoles, TVs, cars, and home appliances. This has caused new challenges for cookies as a mechanism for state management.

As you may recall from the earlier security badge analogy; cookies are extremely specific, and set on a 1:1 basis. That is, the cookie can only be read by the web server who assigned it. In addition to being specific to an individual website or domain, cookies are also specific to (meaning they cannot be shared across): each unique device, each login within a device, and each browser or app connecting to the Internet. None of these cookies can interoperate with other cookies without the help of supplementary technologies like JavaScript or HTTP redirects.

This means when a visitor accesses a weather website via her laptop she is issued a cookie so that the website remembers her location (for the weather forecast). If she then uses her mobile phone to access the same content, the publisher has no way of knowing it's the same visitor—unless she logs into a unique account. This burdens the visitor with having to re-input her location information into each separate device, application or browser she uses to obtain the weather forecast. This also inhibits content and service providers from being able to deliver consistent and customized experiences; including honoring previously set privacy preferences, across this range of digital touch points.

The diagram to the right illustrates how one visitor can potentially generate multiple, disassociated cookies by checking the weather, from a single weather provider, across devices without logging into a unique account.

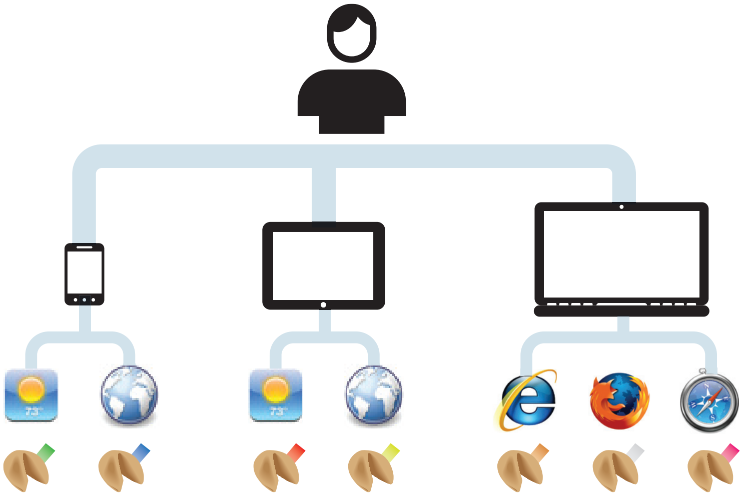

In the case of mobile devices, additional fragmentation may occur within multiple applications on the same device. For example, mobile app publishers and developers may not have access to cookies and will instead often rely on anonymous IDs provided by the operating system (such as the Advertising IDs available on Android or iOS). For mobile web, publishers and content providers may use cookies to support advertising and content delivery functions to some extent, however, some browsers, such as the Safari browser, do not support third party cookies at all. To work around the limitations in both mobile apps and web sites, some publishers and third parties elect to use advanced algorithms to statistically infer an ID using the data available. This leads to even further fragmentation of data, transparency and control for consumers.

Consumers are using a wider array of Internet-connected devices and desire more personalized digital experiences and better control over their digital data and privacy choices. Online publishers are relying on a growing number of third parties to deliver engaging digital experiences and more relevant advertising. The requirements for a persistent and anonymous state management mechanism to achieve these goals have long surpassed the capabilities of the cookie.

6 Addressing the State Management Challenge

The Future of the Cookie working group set out to provide guidelines and criteria for a successful approach to the evolution or replacement of the cookie.

6.1 Stakeholder Groups and Guiding Principles

Prior to researching potential solutions, the working group identified three stakeholder groups with the most pressing needs for effective state management:

- **Consumers** – End users who, via publishers, consume digital content and services
- **Publishers** – Creators, facilitators and/or owners of content (e.g. The New York Times, Yahoo! and Twitter) who provide consumers access to their content and allow third parties (e.g. advertisers) to reach those consumers
- **Third parties** – All companies who deliver content, services or advertising to consumers through publishers

With the understanding that the best solution should meet the needs of all stakeholders, the working group developed a high-level set of guiding principles for each group of stakeholders above. These principles were designed to be used as evaluation criteria against which to consider any and all potential solutions. The working group did not limit evaluation criteria to what was known to be currently available; rather, the working group aimed for a "blue sky" scenario honoring what each stakeholder class ultimately wants.

6.2 Guiding Principles For Consumers

It is absolutely crucial for the consumer's needs and concerns to drive this process. It's their engagement, consumption of content and ultimate spending which drives the economics of this industry. Without full consumer transparency, control and a fair value exchange, publishers and third parties will continue to be limited in their ability to innovate.

Principle	Description
Single privacy dashboard	One place for consumers to see what identities and additional states are stored about them, and where. (Single Place for View-Only)
Universal privacy view	One place for consumers to rea d about them, and where, across all domains and services, via all browsers, devices, logins, and mobile apps. (Minimally a Single Place to say Yes or No to Everyone)
Comprehensive privacy controls	One persistent set of controls for consumers to opt-in, modify, purge or opt-out of data collection or data transfer, in whole or in part, across all parties.
Persistent, universal consumer preferences	Consumer preferences persist across their entire Internet experience, across all domains and services, via all browsers, devices, logins, and mobile apps. It is expected that this requires an authentication mechanism.
Possibility of detecting non-compliant actors.	Simple and easy way for consumers to identify publishers or third parties not in compliance with "the solution".
Free online service	Solution does not require the consumer to install software or hardware, subscribe to a service, pay any money, etc.

Guiding Principles for Publishers / Content Creators

Publishers create the apps, services and content which engage consumers. It is common for publishers to provide free services or content to the consumer in return for the ability to show advertising.

Principle	Description
Single privacy dashboard	One place for publishers to view which third parties are collecting or transferring information what that information is, and how it's being used.
Comprehensive privacy controls	One set of controls for publishers to authorize, limit or deny data collection or data transfer, in aggregate (by party/data type/categories) or in part (for specific parties).
Significantly fewer third party pixels	A reduced number of third party pixels placed within publisher content.
Improved user tracking/targeting	Increased user tracking and targeting capabilities for publishers when compared to traditional cookies.
Reduced cost for privacy compliance	Decreased cost and effort for publishers around end-user privacy matters in general, which might result from browser changes, regulatory requirements, vendor requirements, etc.
Possibility of detecting non-compliant actors.	Easy way for publishers to identify third parties not in compliance with "the solution".
Open competition	Doesn't favor any one specific vendor. Does not provide specific advantages only to certain participants in the industry.
Minimal deployment overhead	Solution does not require publishers to heavily invest in new hardware, software, services or process infrastructure.

6.4 Guiding Principles for Industry Third Parties

Industry third parties are those who aid publishers in the delivery of personalized content, services and advertising. These third parties are instrumental in subsidizing costs to make many Internet services and content free for consumers.

Principle	Description
Decreased segment ramp-up time	Decreased effort and time for third parties to add new tracking partners, segments, sources, etc.
Decreased "cookie churn"	Increased persistence of user identification/tracking.
Lower operating cost	Decreased network traffic and related infrastructure costs for third parties
Better cross-device tracking	More effective user tracking for third parties across devices, logins, browsers, apps, etc.
Better consumer transparency/control	Easy way for third parties to provide end-user transparency and control of user data collection, targeting, and transfer across end-users' entire online experience.
Provides for high integrity frequency capping	Less redundancy in ad exposure for third parties through more effective user-based frequency capping, if the ad server took part.
Less redundant data collection/ transfer	A reduction in the duplication of work collecting user events and transferring this data between third parties.
Reduced regulatory threats	Reduced regulatory threats for all third parties in compliance with "the solution".
Clarifies value to consumer	"The solution" provides for a very clear relationship between targeted online advertising and free content/services.
Vendor agnostic	Open competition: doesn't require commitment to a specific vendor, or provide advantages to only certain participants in the industry.
Minimal deployment overhead	Solution does not require third parties to invest time/effort in new hardware, software, services or process infrastructure.

7 Evaluation of Solution Classes

After achieving consensus on the guiding principles and evaluation criteria, the working group discussed several high-level solution classes. The current cookie approach was used as a baseline for evaluating alternative solution classes. The group reviewed seven distinct approaches to state management—half of which are in-market today and half of which represent new concepts. The goal was not to evaluate individual technologies, but rather high-level solution classes. For example; there are many technologies designed to allow browsers and devices to pass a unique tracking ID (for use in advertising) to third parties. However, as opposed to evaluating specific technologies, the group evaluated the entire solution class against other solution classes in order to be most useful and vendor agnostic. As a result, the seven distinct approaches were combined to create the following five solution classes, roughly based on where and how state is managed:

- **Device-Inferred State** – State managed through the use of IDs inferred using statistical algorithms applied to information passed by the device, browser, app or operating system.
- **Client-Generated State** – State and preferences managed from within the client (such as the browser, app, or operating system) and passed to third parties within the ecosystem. Examples in the market today include the Advertising ID on iOS and Android.
- **Network-Inserted State**– State and preferences managed via IDs set by third party intermediary servers that sit between the end-consumer's device and the publishers' servers. Examples include content distribution networks, Wi-Fi or wireless proxy servers and ISPs. This is a concept not broadly offered in the market today.
- **Server-Issued State** – State and preferences managed via HTTP cookies set between each server domain and browser client, often via Web beacons or pixels. This is the incumbent approach which is ubiquitous in the market today.
- **Cloud-Synchronized State**– State and preferences managed via IDs set by a centralized service that all parties agree to work with. This is a concept not broadly offered in the market today.

7.1 Summary Evaluation

Each solution class is evaluated against each guiding principle as follows:

- 🟢 The state management solution class includes core capabilities that support the principle.
- 🟡 The state management solution class may support the principle in certain implementations.
- 🔴 The state management solution class has significant challenges to overcome in order to support the principle.

	Guiding Principle	Server-Issued Solution	Device-Inferred Solution	Client-Generated Solution	Network-Inserted Solution	Cloud-Synchronized Solution
Consumer	Single privacy dashboard	🔴	🟡	🟡	🟡	🟢
	Universal privacy view.	🔴	🟡	🟡	🟡	🟢
	Comprehensive privacy controls	🟢	🟡	🟢	🟢	🟢
	Persistent, universal consumer preferences	🔴	🟡	🟡	🟡	🟢
	Possibility of detecting non-compliant actors	🔴	🔴	🔴	🔴	🟡
	Free online service	🟢	🟢	🟢	🟢	🟢
Publisher	Single privacy dashboard	🔴	🟡	🟡	🟡	🟢
	Comprehensive privacy controls	🔴	🟡	🟡	🟡	🟢
	Significantly fewer 3rd party pixels	🔴	🟢	🟢	🟢	🟢
	Improved user tracking/targeting	🔴	🟡	🟢	🟢	🟢
	Reduced cost for privacy compliance	🔴	🟡	🟢	🟢	🟢
	Certified participant visibility	🔴	🟡	🟢	🟢	🟢
	Doesn’t tilt towards a specific vendor.	🟢	🟡	🟡	🟡	🟡
	Minimal deployment overhead	🟢	🟡	🟡	🟡	🟡
Industry	Decreased segment ramp-up time	🔴	🟢	🟢	🟢	🟢
	Decreased “cookie churn”	🔴	🟡	🟢	🟢	🟢
	Lower operating cost	🔴	🟡	🟡	🟡	🟡
	Better cross-device tracking	🔴	🟡	🟡	🟡	🟢
	Better consumer transparency/control	🔴	🔴	🟢	🟢	🟢
	Higher integrity frequency capping	🔴	🟢	🟢	🟢	🟢
	Less redundant data collection/transfer	🔴	🟡	🟢	🟢	🟢
	Reduced regulatory threats	🔴	🟡	🟢	🟢	🟢
	Clarifies value to consumer	🟡	🔴	🟡	🟡	🟡
	Doesn’t tilt towards a specific vendor.	🟢	🟡	🟡	🟡	🟡
	Minimal deployment overhead	🟢	🟡	🔴	🔴	🔴

8 Other Important Considerations

In addition to the above guiding principles, other important success criteria were considered.

8.1 Ease of Industry Adoption

Cookies are used ubiquitously today for state management, supporting many features and functionality beyond managing privacy choices and targeted advertising. Due to this ubiquity, mechanisms that are most similar to the cookie/pixel-based status quo will be the easiest for publishers and third parties to adopt. Mechanisms which fundamentally change the way user identifiers are generated and/or stored will require higher risk and more sweeping changes within the industry. Rewriting Internet Standards, for example, could take years to accomplish and is fraught with execution risk. The ideal solution then, is the one that strikes the best balance between meeting the most guiding principles and providing the least effort for the industry to deploy.

8.2 Open Access and Open Competition

Cookie technology is not owned or licensed by one party—rather, it was built into Internet Standards early on and made available to everyone. This allowed the entire Internet industry to embrace cookies for state management without one company having a technology advantage over another. As a result, the industry was able to innovate more rapidly upon a common and open standard, and compete on the basis of value-added features and functionality. The cookie showed us that open access and standardization led to op en competition, which allowed the industry to advance as a whole at a much faster pace.

Several of the solution classes evaluated by the working group resemble current in-market technologies which are owned, controlled and/or licensed by select commercial entities. These entities are ultimately responsible to their shareholders and do not exist to serve the industry or consumers as a public good. Ideally, the right state management solution for the industry—one that properly fuels the next wave of digital innovation—will be openly accessible to, accepted by, and embraced as a standard by all industry participants, just as cookies were. Therefore a certain level of consideration should be given to a model that is owned or controlled by the interest of select for-profit entities. This isn't to imply that the technology must be written into Internet Standards; however, it does mean that a commercial model should be carefully considered and the solution should provide a foundation for earning widespread adoption and trust by the industry and consumers.

8.3 Cross-Platform State Management

Third parties have recently begun to create custom solutions that allow for cross-platform state management—that is, identification of the same visitor across devices, web domains and apps. However, the working group believes that a consistent approach to cross-platform state management is a prerequisite to, and a critical foundation of, advances in cross-platform state management solutions. That is, until a new state management mechanism is adopted by the industry as a standard, replacing the cookie and sufficiently addressing the stakeholder guidelines indicated above, current approaches to cross-platform state management are subject to continued risk of non-standardization, privacy concerns and therefore narrow industry adoption.

Each solution class evaluated by the working group represents a potential path to cross-platform state management—though each has varying degrees of "fitness for purpose". Many third party cross-platform solutions currently in-market most closely resemble the "device-inferred state" solution class. While an evaluation of state management solution classes against the guidelines should be considered first and foremost, consideration should also be given to consumer trust factors pertaining to whether cross-platform state management is opt-in (and explicitly authorized) versus opt-out (and implicitly inferred).

8.4 Fit Within Existing Privacy Programs

While the focus of this paper is technical in nature, it is important to note that privacy advocates and industry trade groups have made tremendous progress with standards and self-regulation pertaining to behavioral targeting, consumer transparency and control, publisher privacy policies, and other policy concerns. The best solution must extend the excellent work done by the DAA (aboutads.info), the Network Advertising Initiative (networkadversting.org), and other organizations.

8.5 State Management and State Synchronization

While most of the solution classes focus on creating state; the cloud solution class provides a mechanism for bridging other solution classes, thus enabling additional features.

9 Summary

The realities of the evolving digital ecosystem have resulted in the cookie being pushed beyond its useful and intended purposes. This multi-device, multi-platform, multi-environment reality has presented new challenges that the cookie is not able to address. What began as a simple state management solution has become the foundation of a complex and valuable online marketplace. This marketplace has grown to include thousands of stakeholder companies and digital consumption now extends across smartphones, tablets, TVs and an ever-evolving array of Internet-enabled devices. In this new reality, the cookie cannot serve as the foundation for the next generation marketplace. New approaches are required.

Several things became clear during this state management solution class analysis. First, there is no perfect solution available today. In fact, most solutions that are viable today have significant limitations and the solution class which satisfies the most guiding principles (Cloud based) is, at best, very nascent concept in the market. Second, the reality of today's market is that several companies have developed consumer footprints large enough to (arguably) make a proprietary solution to the state management challenge a reasonable possibility. While the adoption of a proprietary solution as an industry standard would serve the multi-device consumer, it would also fundamentally change the current state management landscape, which is based on an open standard. The guidelines and analysis here are meant to provide a framework for the ongoing evaluation of this challenge and the inevitable tradeoffs that will have to be made.

10 Recommended Next Steps

Next Steps Recommended:

- Publish Best Practices for Implementation for each of the solution classes
- Observe the industry adoption of these technologies

11 Appendix A — Examining the Solution Classes in Depth

11.1 Server-Issued State (Cookies)

11.1.1 Description

This is the incumbent solution class which uses cookies to store information that publishers or third parties can access when requests are made of their servers. In the case of third parties, cookies are usually set via a process where web beacons call out to third party servers requesting 1x1 pixels to be placed on publisher pages. Note that this solution class may also include other mechanisms for storing information within the browser client, such as HTML5 local storage objects.

A cookie (also known as an HTTP cookie, web cookie, or browser cookie) is a small piece of data sent from any Web server and stored in a visitor's web browser while the visitor is browsing that website. Every time the visitor visits the website, the browser sends the cookie back to the Web server so that the website can remember who the visitor is. Cookies were designed to be a reliable state management mechanism for websites to remember information (such as a user ID or items in a shopping cart).

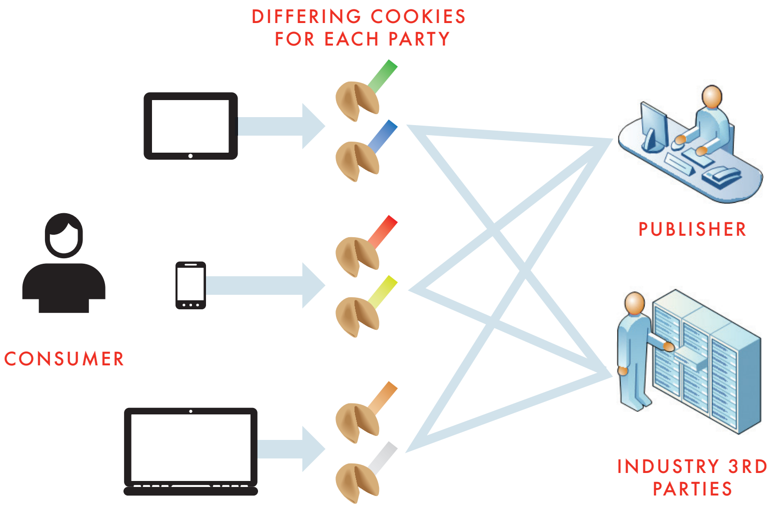

The core challenge with cookie technology is that it's highly decentralized. Cookies can only be set on a 1:1 basis. That is, the cookie can only be read by the web server who assigned it. In addition to being specific to an individual website or domain, cookies are also specific to (meaning they cannot be shared across): each unique device, each login within a device, and each browser or app connecting to the internet. None of these cookies can interoperate with other cookies without the help of supplementary technologies like JavaScript or HTTP redirects. For security reasons, browsers will only send cookies back to the originating servers.

As a result, another core challenge with cookie technology is that data sharing is not easy or efficient for any party. The industry has addressed this with Web beacons and "redirects" that, when placed on publishers' pages, allow for information to be passed to third party recipients and for those third party ad servers to set cookies of their own within the visitor's browser. The entire ad industry relies on data exchange, and all data exchange relies on these mechanisms, which must be duplicated for every party and every user. That necessary duplication is problematic to consumers, publishers and industry third parties.

11.1.2 For Consumers

Cookie technology, as an HTTP standard, was designed for one purpose: state management between a Web browser and a Web server. Since there has been no standardization for how cookies should be used many parties use them in different ways, for different purposes—all of which are opaque. This makes a single privacy dashboard for consumers extremely difficult to offer. It can be difficult to know what the information stored client-side (within hundreds of cookies) actually means, and it may pale in comparison to the proprietary information fragmented across hundreds of different servers.

Furthermore, the mechanisms necessary for data exchange between publishers and third parties result in a massive number of Web beacons placed on publishers' pages, which increases latency and negatively affects the user experience.

For all of the above reasons and more, concerned consumers often erase their cookies (either manually, or using third party security software that does it automatically) or disallow the setting of cookies altogether. The resulting instability of cookie availability, combined with the fact that cookie support varies by device, operating system, app and browser, results in cookies being an unreliable mechanism to provide consumers with comprehensive and persistent privacy preferences. The lack of stability and persistence with cookies also makes it difficult for consumers to tie their devices together, in order to apply privacy preferences universally and consistently across their digital experience.

11.1.3 For Publishers

Within this solution class, the mechanisms necessary for data exchange between publishers and their third party partners result in a massive number of Web beacons placed on publishers' pages, which increases latency and negatively affects the end-user experience.

With parties using cookies in opaque ways for varied purposes, a single privacy dashboard for publishers is extremely difficult to achieve. Publishers see an exponential number of third party requests on their pages, though within this solution class it is nearly impossible for publishers to know who is collecting or transferring proprietary information on their pages, what that information is, and how it's being used. Thus they are unable to go one step farther and authorize or deny the data collection.

Since cookies are increasingly unreliable as a mechanism for state management, they are increasingly unreliable as a mechanism for offering personalized content, features, and advertising—especially if the publisher depends on partnerships with third parties to achieve these things.

As cookies have decreased in reliability, parties have increased the volume of Web beacons in an attempt to compensate. This has escalated privacy concerns among consumers and regulators, and has also increased costs for publishers to remain compliant with privacy laws and best practices.

As cookie technology is an open HTTP standard not owned or licensed by any single party (or set of parties), the use of cookies does not inherently advantage one publisher over another. Absent any industry policies treating some cookies differently than others, this solution class involves a level playing field, which opens up competition and therefore innovation.

11.1.4 For Industry 3rd Parties

Within this solution class, third parties must rely on Web beacons or pixels placed on publishers' pages to build audience segments across the Web, which means a particular user can't be included in an audience segment unless they've reached a page with that third party's beacon on it. It's not possible for every third party to reach each Internet visitor every day, so it can take weeks to build any given segment. Furthermore, for every change to an audience segment, or addition of a new partner, new Web beacons must be deployed and those audience members need to be seen again. This leads to massively redundant and voluminous deployment of Web beacons.

With an ever-increasing number of third party Web beacons being deployed, concerned consumers often erase their cookies (either manually, or using third party security software that does it automatically) or disallow the setting of cookies altogether. That's why this solution class is characterized by "cookie churn", where industry third parties experience rapidly diminishing audience segments until they can find that visitor and reset the cookie—again and again.

The lack of stability and persistence with cookies, and the fact that they're set on a 1:1 basis, makes it difficult for third parties to target audiences across devices, apply persistent privacy preferences, frequency caps or targeted advertising consistently across consumers' entire digital experience. For all these reasons, third parties will experience increasing operating costs under this solution class as they continue to deploy more Web beacons and experience diminishing returns.

Cookie technology, because it's an open HTTP standard not owned or licensed by any single party (or set of parties) and used ubiquitously today, does not inherently advantage one third party over another. Absent any industry policies treating some cookies differently than others, this solution class involves a level playing field, which opens up competition and therefore innovation. However, if industry policies treated third parties differently than first parties, in terms of "allowable use of cookies", that could materially tilt the playing field towards the few large first parties that don't depend on third parties to deliver personalized features, services and advertising.

11.2 Device-Inferred State

11.2.1 Description

This solution class is characterized by the use of statistical algorithms that, using information passed by the device, browser, app or operating system, infer a user ID which can then be used by publishers or third parties to manage state.

Statistical identification of devices works across multiple apps or programs on a single device. It establishes state that is tied to a specific device, which is accessible to multiple third parties provided they are using the same algorithm.

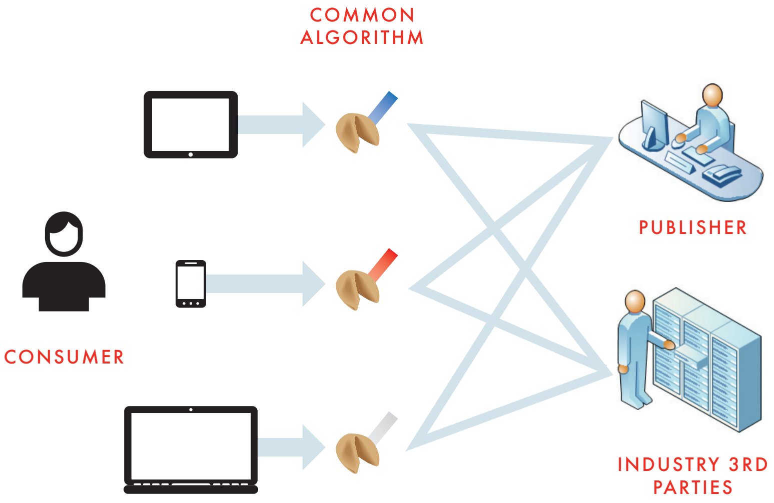

The core challenge with this solution class is that the IDs utilized for state management are based on probability, which means there will always be a level of inaccuracy and instability. The same ID could apply to multiple devices and/or consumers - especially in scenarios where an IT department sets up devices in a standard way for a corporation. Furthermore, IDs could change frequently for the same consumer due to network or IP address changes, device or O/S updates, change to browser add-ons, etc. As a result, a truly stable and persistent ID, necessary for any reason from maintaining a shopping cart or honoring consumer privacy preferences, would be nearly impossible to maintain. For this reason, statistical identification solutions should be deployed in combination with a deterministic system for honoring explicit consumer preferences, including opt-out decisions. For example, a company might rely on statistical identification for tracking conversions, while relying on a consumer opt-out mechanism based on cookies (server-issued state) or IDFA (client-issued state).

11.2.3 For Consumers

Since an inferred, statistically determined ID would always result in a level of inaccuracy and instability, and be specific to the device, a single privacy dashboard for consumers that relies entirely on statistical IDs would be extremely difficult to offer with any persistence or integrity. It would also be difficult for consumers to tie their devices together, in order to apply their privacy preferences universally and consistently across their digital experience. Coupling the statistical solution with other solution classes and connecting to a central source (such as a cloud-based solution) would be necessary to provide a single privacy dashboard for consumers, and allow them to persist those preferences across multiple tied devices.

There may be an opportunity for standardization of a statistical algorithm that is shared industry-wide. Under that scenario, a centralized entity may be able to offer a consumer privacy dashboard if all participating parties provided information. This might also facilitate a central list of companies so that consumers had visibility to what companies don't comply, though there still would be limited transparency for consumers during the time state management is actually being executed (by publishers or third parties).

11.2.4 For Publishers

This solution class would not independently offer publishers a single source of transparency and control over the parties collecting data within their content. However, the solution could be coupled with a cloud-based solution to allow publishers to narrowly restrict any/all third party requests. One consideration with this solution class is that, unless controls were offered to and supported by publishers, the existence of a global ID may actually enable third parties to get access to publisher audience data without publishers being aware, or transfer data between each other without the knowledge of publishers. To address this concern, a device-based implementation could enable publishers to check a centralized list to know which third parties are supporting data transfer standards, keeping in mind that there would still be difficulty validating that the third parties are truly in compliance. On the plus side, this solution class makes it easier for publishers to work with their direct data partners or a third party data management platform.

Assuming this solution class would be deployed as a standard and shared industry-wide, it obviates the need for multiple pixels on publisher pages that serve the sole purpose of mapping proprietary IDs between third parties. Elimination of the need for "ID mapping" would result in decreased latency, improved page load times and an improved consumer and advertiser experience.

Since an inferred, statistically determined ID would always result in a level of inaccuracy and instability, some publishers reliant on the mechanism to work with third parties may not find the economics to be in their favor. Particularly given a proprietary approach by a third party vendor, the costs associated with the approach would certainly exceed the cost of cookies (which incur no licensing fees).

Lastly, any publishers with proprietary infrastructure, features or services may have to retool, from cookie-based storage to either a cloud or server-based data store which is tied to the statistically determined ID and actionable at run-time. These publishers might also find an increase in vendor costs as their providers adapt to this methodology.

11.2.4 For Industry Third Parties

Provided that the solution class were deployed as an industry-wide standard, and in a manner that mitigated the ID stability concern for certain use cases, the benefits to industry third parties are several-fold. First, operating costs for third parties would be reduced by eliminating the need for multiple pixel placements that serve the sole purpose of mapping proprietary IDs between industry third parties. Second, a single universal ID results in less redundant data collection and would allow for more rapid development, deployment, and transfer of proprietary audience segments. Lastly, third parties would be able to support higher integrity frequency capping of advertisements, which in turn would reduce impression waste.

This solution coupled with a cloud-based solution may also give industry third parties the benefit of better cross-device tracking and reduced regulatory threats; as centralization would be helpful in respecting consumer preferences. That said, since this solution class is not inherently visible or easy to manage for consumers, it does require a heavy focus on consumer education

An inferred, statistically determined ID always results in a level of inaccuracy and instability and as such, some third parties may find the lack of accuracy to be unacceptable. Additionally, the economics may not be favorable if the solution relies on proprietary technology by a specific vendor (compared to cookies, which do not require licensing fees). Generally speaking, third parties would see more benefits the more broadly standardized a solution of this class were. While open competition can be good for industry, in this case the benefits of a single universal ID can only result from a standard, or a concentration of power within a monopoly or oligopoly.

11.3 Client-Generated State

11.3.1 Description

This solution class manages state and preferences from within the device or client (i.e. the browser, app, or operating system) and passes a static ID on any call to third parties within the ecosystem. Examples in the market today include Apple's AdID, and similar developments by Google and Microsoft.

Client-side code works across multiple apps and programs on a single device or operating system. It establishes a persistent state that is tied to a specific device. This state then becomes accessible to multiple servers in different domains. Depending on implementation, this solution class may provide identification to only specific clients on a device (i.e. Apple ID only available through native applications and not through the mobile browser or Safari).

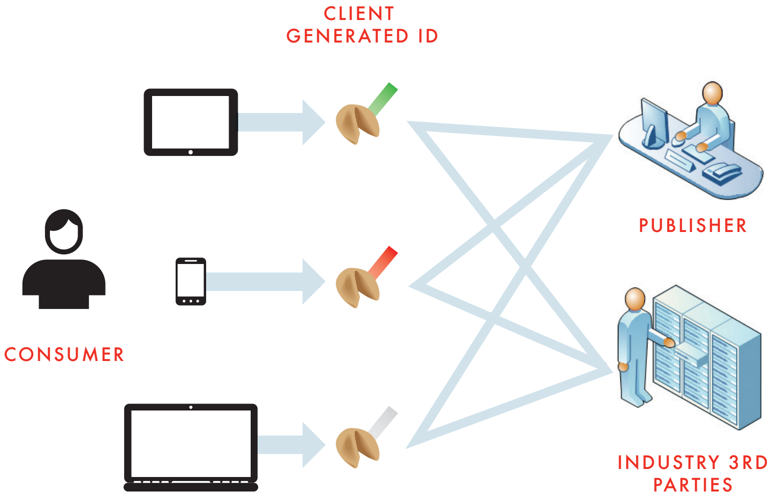

The main drawback to this solution class is its device-centric or client-centric nature. State management across devices, used to target ads to or honor a consumer's preferences across a number of their Internet-connected devices, would require a centralized login (such as a cloud-based solution).

11.3.2 For Consumers

Because of the device-centric nature of this solution class, a single privacy dashboard for consumers would be extremely difficult to offer with any persistence or integrity. Consumers would only be able to tie their devices together by providing additional information to a central source (such as a cloud-based solution) which would map the data of login information across devices in order to apply privacy preferences universally and consistently across the consumers' complete digital experience. Alone; this solution would never provide a global view of state and preferences across all devices.

With this solution each creator of client-side IDs, whether browser, operating system, or application, should be able to recognize non-participants within their own infrastructure as well as those deliberately trying to circumvent permissions or privacy. This information could then be propagated back to the consumer.

A big benefit to this solution class is that privacy features can be more clearly surfaced to consumers within the client they're using, at the time of use. There is also the opportunity for privacy features to be more elaborately designed and better enforced by each client—though lack of standardization or centralization around privacy features could confuse consumers.

11.3.3 For Publishers

To create a first-party cross-device dashboard, publishers would need to join server logs with their own user login info; however they would still lack direct details about the third party data sharing that occurs outside of their owned and operated sites. As written, the solution class does not specifically address this, but an execution of this solution class could be extended to do so.

A benefit to the publisher is the reduction in number of pixel calls. Depending on implementation, this proposal negates the need for pixels entirely—if a publisher is willing to utilize a server-to-server based data sharing implementation. However, due to the proliferation of third-party data sharing in the industry a publisher may not have insight into the data shared outside its purview.

Overall, having client-generated code that is sent on all requests creates standards for state management and sharing that are currently unavailable in the cookie-based model. By virtue of having standards, it is expected that there will overall be less churn in IDs and, more importantly, a reduction in cost for publishers to adopt and maintain this scheme.

For the purposes of introducing this proposal class to the working group the description is patterned heavily after IDFA, Apple's client-generated code solution. Adherence to that model is not a requirement; the ID generation can be technology agnostic as long as the ID is shared.

As mentioned above, adoption is made easier by virtue of having standards around data sharing but there is still overhead in widespread publisher adoption. For browser-based publishers, additional work is needed to identify the appropriate mechanism for inserting the client-side ID and passing it along as necessary. Application developers should find this trivial in terms of development, but rolling out the changes to all their consumers can take a long time depending on user update frequency

11.3.4 For Industry 3rd Parties

Similar to publishers, industry third parties will benefit by having a more standard approach to data sharing. This leads to easier industry adoption and fewer specialized integrations. Currently, user matching and data sharing agreements can require one-off integration types with customized APIs. At least in the short term, this increase in standards will reduce cookie and user churn.

Lower churn and more stability of data require less overhead from industry third parties. While a new identifier in request headers could increase the overall payload per request, the overall volume of data sharing calls among third parties will decrease. Many third parties maintain their own data stores for collection and evaluation of user data and trends; these may require additional resources because of the addition of more data in a header. This may lead to a slight increase in costs associated with data sharing, but it may be an insignificant amount depending on the implementation.

Cross-device tracking would depend on how well the system can get this information and enforce controls and is highly dependent on implementation details.

A consistent ID attached to user-set and all-party respected preferences is a strong privacy move by the industry, but changes in the legislative landscape may continue to challenge any form of client identification. This could lead to additional churn from users deleting client-generated IDs on a regular basis, resulting in many of the same overhead issues the industry faces today.

11.4 Network-Inserted State

11.4.1 Description

This solution class manages state and preferences via IDs set by third party intermediary servers that sit between the end-consumer's device and the publishers' servers. Examples include content distribution networks, Wi-Fi or wireless proxy servers, and ISPs. This is a concept not broadly offered in the market today.

Network-inserted stated works across multiple apps/programs on a single device, and can also be implemented to work across multiple devices in the same household. Specifically, a network-based solution could connect devices to a single user identifier if the devices were all registered with the Network (e.g. an ISP) by the same user or household. Additionally, with certain implementations or partnerships in place, a network-inserted ID could work across Networks.

In general, this solution class establishes a persistent state against a single device or household, and makes this available to multiple servers in different domains. The core benefit to this solution class is the reliability of the ID because it is consistently set per the Network's relationship with the consumer.

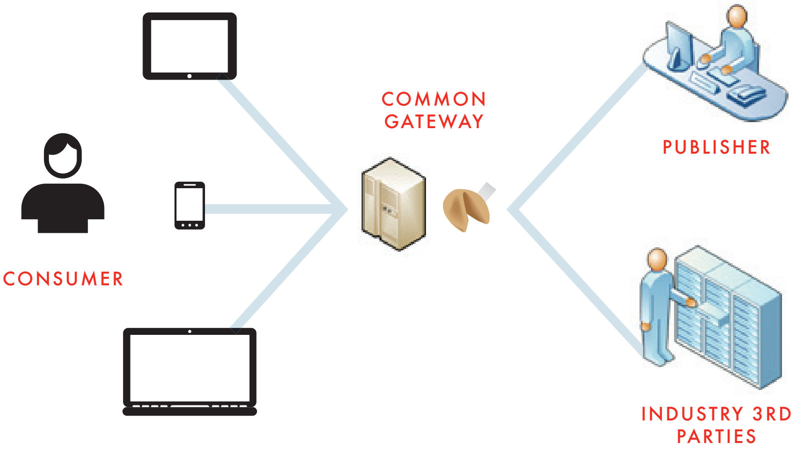

11.4.2 For Consumers

With a per-device execution of this solution class, a single privacy dashboard for consumers, similar to other solutions, would be extremely difficult to offer with any persistence or integrity. Consumers would however be able to tie their devices together by providing additional information to a central source (such as the ISP or Carrier, via a cloud-based solution) that handled mapping of Network-inserted IDs across devices, in order to apply their privacy preferences universally and consistently across their digital experience. Given the relationship with the consumer, the Network may be able to easily provide this option and deliver a global view of state and preferences across all devices.

Because this solution class is built on the existing relationship between the consumer and their network service provider, there is a level of trust inherent in the relationship. Consumers would have the expectation, and service providers would be motivated to meet the expectation, of ensuring that only trusted participants have access to user state. Central availability of information related to participants in the solution as well as non-participants, could easily be propagated back to consumers. Establishing this level of trust would support turning the tide of concerns associated with consumer's discomfort with existing state management solutions, and the resulting reactions and concerns in the regulatory space.

To maintain the integrity of this solution class, Network providers would need to address the potential that identification could somehow be passed on to non-trusted entities. Additionally, a mechanism for monitoring trusted parties to ensure they are upholding best practices would be necessary to satisfy consumer needs for a reliable solution.

11.4.3 For Publishers

As written, this solution class does not specifically address the guiding principle of providing publishers with direct details about third party data sharing that's outside of their owned and operated sites. That said, it is possible that an execution of this solution class may be extended to do so.

A clear benefit to the publisher with this solution is the potential to reduce pixel calls. Depending on implementation, this proposal could negate the need for pixels entirely, if a publisher is willing to change to a server-to-server based data sharing implementation. Additionally, this solution introduces standards for state management that would be based on the consumer's trust with the Network. With the addition of standards and a transparent, choice-focused model, it could be expected that publishers will have a decline in challenges associated with cookie churn and deletion.

Clearly, adoption is made easier by virtue of having standards around data sharing, but, like other solutions, there is still overhead in widespread publisher adoption. Depending on the execution, the publisher/application developer, and their level of experience working with different IDs, the level of impact may vary.

11.4.4 For Industry 3rd Parties

Assuming this solution class would be deployed as a standard and shared industry-wide; there are multiple benefits for industry third parties. First, by eliminating the need for multiple pixel placements that serve the sole purpose of mapping proprietary IDs between industry third parties, it reduces third players' operating costs. Second, a single universal ID results in less redundant data collection, and would allow for more rapid development, deployment and transfer of proprietary audience segments. Third, third parties would benefit from higher integrity frequency capping of advertisements, reducing impression waste.

Coupled with a cloud-based solution, industry third parties may also benefit from better cross-device tracking and reduced regulatory threats, since centralization would be helpful in respecting consumer preferences. Third parties would benefit from the trusted relationship between Network providers and their consumers which is the foundation of the solution class.

With the centralized ID solution, third parties can trade data against a single ID which minimizes the effort of creating unique IDs and data calls for new tracking partners. Additionally, the ID would be omni-present (assuming the execution is applied in a way that extends to Wi-Fi) and would not be subject to situations akin to cookie deletions or inoperability. While the long term advantage is clear, it is important to note that third parties with proprietary infrastructure, features or services may have to retool, from cookie-based storage to either a cloud or server-based data store which is tied to the Network-Inserted ID and actionable at run-time.

The network-inserted ID concept is intended to minimize operational costs by minimizing or even eliminating the need to support other technologies such as cookies. Because in mobile, parties are accustomed to managing multiple IDs, it is likely not difficult to execute; however, some level of coding/infrastructure will be needed to support a new ID solution (which is likely the case with all the solution cases). Depending on the execution, the particular third party, and their level of experience working with different IDs, the level of impact may vary.

11.5 Cloud-Synchronized State

11.5.1 Description

This solution class provides a layer of synchronization atop other solution classes discussed above. This represents state and preferences managed via IDs set and synchronized through a centralized service that all parties agree to work with. This is a concept not broadly offered in the market today.

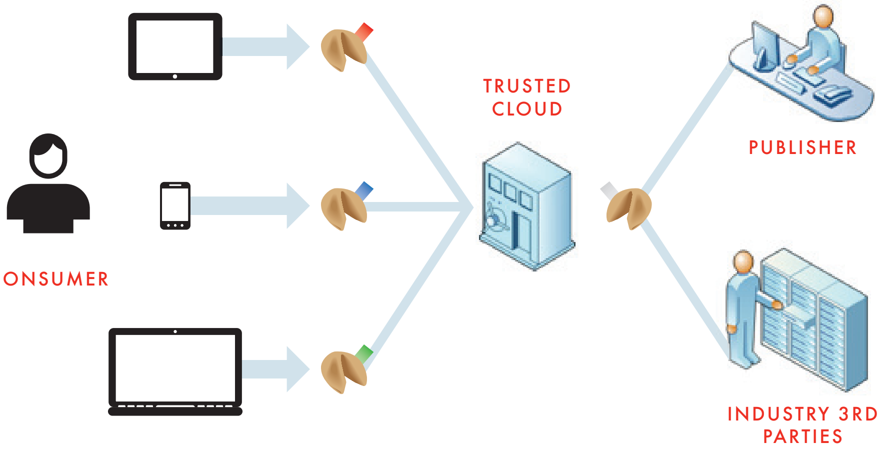

Cloud-synchronized state works across programs and devices. It establishes a persistent state that can be the aggregation of multiple other states (e.g. from other solution classes discussed in this document), and can make this state available to multiple servers in different domains.

This is not a standalone state generation technology, but rather a technology that would allow greater consistency of experience based on user preference and the strengths of the other solution classes.

11.5.2 For Consumers

Establishing a central repository for state, preference, and data permissions management provides opportunities for improving on the current industry preference options, allowing more fine-grained control, visibility, and greater preference longevity.

The cloud-synchronized state could allow users to declare associations between devices, and their IDs from different, previously discussed, solution classes. As such, the cloud-synchronization becomes the glue linking preferences across platforms and devices, solving current inconsistent user experience across programs, apps, and devices.

11.5.3 For Publishers

The centralized repository for state, preference, and data permissions management improves the publisher experience in several core ways:

- Reduction of tags needed for ID synchronization, resulting in a decrease in page load time (PLT), and a reduced exposure to data leakage situations.
- Reduced spin-up time for new data partners, and elimination of retagging to add and remove data partners.
- Better persistence and access to user preferences. User ad-related Opt-Out preferences are often only sent to ad servers and not publishers, which impacts the value of an impression in an unpredictable manner for the publisher.

A cloud-based state synchronization technology is most compelling when it is managed by a neutral body—it must be trusted by companies that would otherwise be competitors. Establishing a singular trusted tracking provider and clearinghouse could result in improvements to the usage of all other solution classes, as the guidelines established by this one party would guide the whole industry.

It is important to note that there are risks of having only one organization through which identity synchronization occurs, and that the oversight of such an organization must be thorough.

11.5.4 For Industry 3rd Parties

A centralized repository for state has the potential to reduce primary infrastructure requirements for industry third parties, by centralizing all beacon calls to one exchange, and to remove the technical requirement of retagging when establishing new business relationships. Therefore, it is expected that such a solution would lower operating costs and barriers to participation for smaller organizations.

The centralization comes at a risk as identified above—while a single organization for synchronization is the most technically efficient, such centralization must be transparent to all participants in order to allay concerns of skew towards any other participants. This market-driven need to be transparent is a strong positive message to consumers and regulators.

The ability of a cloud-synchronized state solution to associate multiple types of IDs with one another allows the strengths of previously discussed state management solution to complement each other, providing opportunities for cross-device tracking, reducing redundant ad exposure, data collection, and data transfer. To do this effectively, the solution would need to solicit user input, thereby strongly identifying value to the consumer.

These upsides do come at the cost of significant changes to the current infrastructure for the industry.

12 Appendix B – Glossary of Terms

Consumers – For the purpose of this document consumers (also referred to as users, end users, and visitors) are end users of Web content and services.

Cookie – A cookie is a small text file sent from a website and stored on a visitor's web browser while the visitor is viewing the website. Cookies can only be read by the assigning website; i.e. websites can't read cookies from other websites.

> **First Party Cookie** – Cookies which are assigned in and by the domain of the website shown in the browser's address bar.
>
> **Third Party Cookie** – Cookies which are assigned in and by a domain different from the website shown in the browser's address bar. These cookies originate from parties who serve content into the webpage you're visiting (e.g. advertisers, plugins and other content providers).

Cookie churn – Cookie churn refers to instances when cookies are deleted or expire. When a cookie is deleted or expires, all associated information is lost (including preferences and state).

Data leakage – The unauthorized or unintentional use of a publisher's consumer's data by an advertiser or other third party outside of the publisher's domain.

Domain – The unique name that identifies a website.

End User – see "consumer" above.

Frequency Capping – The process of restricting the number of times a set of creative or content is delivered to a consumer. The effectiveness of frequency capping is impacted by the efficacy of the state management solution being used.

HTTP cookie – see "cookie" above.

Internet Standards – The collection of specifications that govern the technical execution of communication on the Internet; including HTML, HTTP, and TCP/IP.

Publishers – Creators of apps, web services and content (e.g. The New York Times, Yahoo!, Twitter and others).

State – Persistent attributes associated with a device, web browser, app, household, or other proxy for the consumer. These attributes can include a user ID, opt-out preferences, authentication tokens, and IDs for ad targeting.

State Management – The process of providing the information necessary for first and third parties to deliver personalized information and services to end consumers, and/or respect their preferences for privacy, information transparency and control.

Web beacon – Also known as pixel, a web beacon is a tiny image (usually measuring 1x1) referenced by a line of HTML or JavaScript code embedded into a website or third party ad server to track activity.

Web domain – see "domain" above.

Webpage – A collection of first and third party resources (images, scripts, audio, video, etc.) that are presented through a web browser in a single, unified experience.

Website – A collection of webpages, generally served from the same domain or publisher.