AppNexus

HEADER BIDDING:

THE NEXT EVOLUTION

WHAT EVERY PUBLISHER NEEDS TO KNOW ABOUT DEMAND SOURCES, SERVER-TO-SERVER, AND EMERGING FORMATS

INTRODUCTION

When header bidding started to gain momentum in 2015, it promised publishers the ability to sell every impression at its true market value.

Two years later, it's clear that header bidding has delivered on that promise and had a profound impact on the digital marketplace. Right from the start, a wide range of publishers began seeing greater yields and enjoying unprecedented insight into what their inventory was really worth. Over time, technology vendors and publishers worked together to refine the process further, building new tools that made header bidding faster, more effective, and easier to manage. With each improvement, publishers got better results, and new companies decided to give the solution a shot. As a result, the vast majority of today's top publishers use header bidding.

But for as much as header bidding has grown over the past two years, there is yet more change on the horizon. Indeed, when we last **explored header bidding technology** in August of 2016, the biggest questions facing publishers revolved around which wrapper they should use to manage their auctions. Since then, the conversation has evolved to encompass new priorities, new technologies, and of course, new questions. Today, publishers are increasingly focused on finding the right demand sources, pushing into new formats, and looking for ways to improve header bidding's impact on user experience.

This whitepaper aims to help publishers navigate this complex ecosystem by addressing their most pressing concerns. Ultimately, we hope to give you the information you need to maximize your header bidding returns, both now and in the future. Here's a taste of what we'll be covering:

The Standardization Of The Wrapper And A Renewed Focus On Demand

We've been helping publishers optimize their header bidding setups from the beginning, and we know what a wrapper needs to do to make them successful. We can tell you exactly what standards your header bidding container needs to meet so you can focus your attention on building the best possible list of demand partners for your inventory.

The Great Debate: Server-to-Server Header Bidding Vs. Client-Side

Server-to-server (S2S) header bidding promises to reduce latency and expand the number of partners publishers can work with. However, this setup does come with a few potential drawbacks. We'll explain the key differences between S2S and client-side header bidding, so that you can decide which solution -- or combination of the two solutions -- makes the most sense for you.

The Expansion Of Heading Bidding Into New Territory

Until recently, header bidding was almost exclusively used for selling web display inventory. But now, technology advances are making it possible to bring header bidding to the mobile app, video, and native environments. We'll discuss how header bidding can help publishers get the most out of these formats.

But before we dive into the future of header bidding, let's take a quick look back at the past. After all, it's hard to know where something is going if you don't know where it came from.

PRELUDE: HEADER BIDDING 101

The waterfall was leaking money with every new auction. Publishers needed something better.

Given header bidding's massive popularity, it can sometimes be difficult to remember that it's still a relatively new technology. In fact, it was just a few short years ago that the monetization technique was first introduced to help publishers allocate their inventory more efficiently.

Header bidding was created to address a big pain point in programmatic advertising: Publisher ad servers have historically been unable to efficiently evaluate bids from multiple demand partners at once. Instead, an ad server would offer an impression to the publisher's various demand partners one by one, awarding the impression to the first that met its decisioning criteria – usually based on a mix of factors like revenue, campaign targeting preferences, and how often or recently the user had seen the ad in question. This worked fairly well when publishers were only running direct campaigns, as it was easy for them to prioritize campaigns based on the static CPMs they had agreed to in advance. But once publishers started adding programmatic partners that bid a different CPM for each impression, the ad server became woefully outdated. It had no good way of prioritizing these new demand partners without knowing in advance what CPM value they would bring.

To make up for this shortcoming, publisher ad servers were built with decisioning logic commonly referred to as "waterfalling." In a waterfall setup, publishers estimate the revenue per impression they'll get from each of their programmatic partners based on how much they've bid in the past. That static value determines the order in which those partners are offered impressions. But since programmatic buyers produce a different bid for every impression, there's no way of knowing whether the first partner in the waterfall is actually the one prepared to offer the highest bid. In many instances, this has caused publishers to sell impressions at a lower price than what a partner further down the waterfall would have been willing to pay. Every time this happens, publishers lose revenue that should have been theirs.

Header bidding creates a level playing field between demand sources, allowing publishers to sell every impression to the right bidder.

These problems were compounded by the fact that many publishers were using ad servers owned by companies that also ran programmatic exchanges. By favoring their own demand with last-look privileges and opaque auction logic publishers couldn't see for themselves, these vendors might have given their own exchanges an unfair advantage -- many in the industry were certainly suspicious. This limited competition meant lower CPMs for publishers and even prompted some media buyers to stop participating in auctions altogether due to their low chances of winning, depressing CPMs even further.

Fortunately, it wasn't long before header bidding came to the rescue.

Header bidding is a monetization solution that allows publishers to source bids from all of their demand sources at the same time by making ad calls to all of them simultaneously. This eliminates the inefficient hierarchy of the waterfall, ensures that every demand partner gets a chance to bid on every impression, and increases yield for publishers by allowing greater competition.

After those simultaneous ad calls, the header bidding process varies depending on what solution the publisher is using. Many header bidding solutions then hold a first-price auction in the header of the page and then send the winning bid to the publisher's ad server. Other solutions, including Prebid, don't hold an auction at all, but instead simply pass all bids through to the ad server. We designed Prebid to do this because we believe it's more efficient for all submitted bids to be subject to the publisher's yield management logic in the ad server, which is the only place where a fully unified auction can take place consistently. **We've found** that when wrappers try to handle things like price floors and bid reduction, it skews the competition and results in lower yield for publishers.

Regardless of where the auction takes place, the ad server then has to decide whether it should award the impression to the winning header bidder, to the ad server's programmatic demand, or fill it with creative from an existing direct campaign.

Header bidding lets publishers ditch the waterfall and have all their demand partners -- both direct and programmatic -- compete against one another on even footing and in real time. That doesn't just mean higher prices for publishers. It also gives more buyers access to better inventory they'd never have had a shot at before, which makes their campaigns more effective.

Despite header bidding's early success, we've yet to realize its full potential.

Beyond that, header bidding kicks off a virtuous cycle that improves the digital media landscape for all. The heightened competition increases overall CPMs for publishers, giving them the money they need to hire journalists and developers who help them publish great content. Users come back for more of that content, which drives better results for advertisers and prompts them to continue increasing their spend. Everybody wins.

Of course, for all that header bidding has brought to publishers and advertisers these past two years, it wasn't quite perfect right out of the gate. As pubishers expanded their header bidding auctions, they noticed that adding too many demand partners can cause latency and place an unreasonable strain on the ad ops team charged with integrating them. And since header bidding grew out of display-specific needs, there wasn't yet a way to bring its benefits to the full array of publisher inventory like video, native, and mobile.

Over the course of this whitepaper, we'll explore three trends that show how header bidding technology is evolving to meet these challenges. They are: the maturation of the header bidding wrapper, the rising popularity of server-to-server header bidding, and the expansion of header bidding into the emerging formats of video, mobile app, and native advertising.

HOW HEADER BIDDING INTEGRATIONS WORK

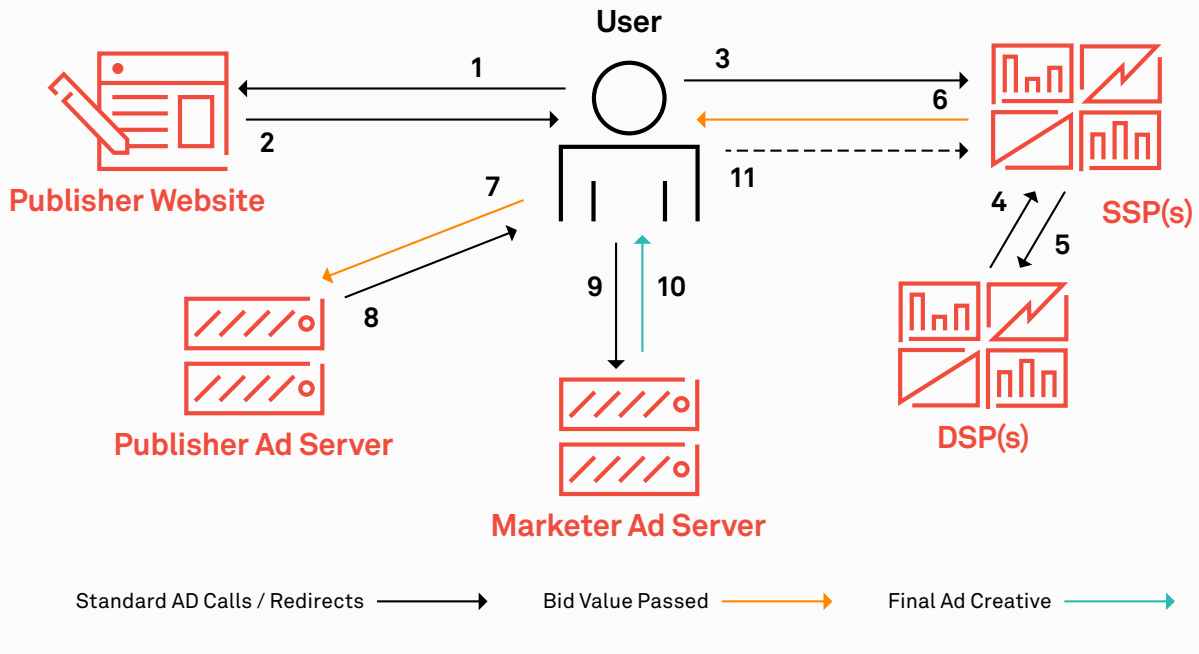

CHAPTER 1: WRAPPERS AND DEMAND

The wrapper makes it easier for publishers to manage their header bidding partners.

While publishers were initially thrilled with header bidding's impact on their CPMs and yields, many of them became frustrated by how long it took them to add new partners to the auction. Back then, every additional demand source required publishers to endure the time-consuming process of adding custom javascript to their site's header to both call and digest responses from the new partner.

In order to ease the pain, advertising technology companies introduced a tool called the **header bidding container, otherwise known as a wrapper**. The wrapper is a single piece of code that holds the basic logic for every one of a publisher's header bidding partners. After one simple implementation, publishers can easily add or remove partners from the header bidding auction, allowing them to scale their header bidding solution much more efficiently.

At AppNexus, we quickly came to the conclusion that wrapper technology would be essential to making publishers successful with header bidding. We also wanted to make sure that any publisher could take advantage of header bidding, and do it in a simple, safe, and easy-to-use way. That's why we decided to build and release a free, open-source, open-source header bidding header bidding technology: **Prebid**.

No matter what wrapper you choose, what's really important is that you're using it to connect with as many quality demand sources as possible.

Our biggest goal with Prebid was to create a tool that would be easy to implement and compatible with any demand source. Going open source has been a huge help in meeting that goal. We have an invaluable community of contributors who maintain Prebid's code and build new features. Plus, an open-source adapter makes it easy for any demand partner to build their own adapter. Right now, there are around 75 adapters, which gives publishers plenty of ways to connect to the buy-side.

But regardless of which wrapper you choose, what's most important is that you use it to connect to a broad group of quality demand partners who make sense for your inventory. That's the key to building a robust, liquid marketplace for digital ads -- one that allows buyers and sellers can make frictionless, fairly-priced transactions that put the right ads in front of the right users.

So, how do you actually go about building a strong demand profile? It's not a matter of just adding as many partners as you can to your header. You need to evaluate your header bidding partners on four crucial criteria.

The first is **wide market access**. By this, we mean that all your demand sources need to bring different advertisers to your inventory. Otherwise, you'll just have the same buyers bidding on your impressions through different exchanges.

There are four keys to finding the right header bidding demand mix: wide market access, quality, high performance, and open competition.

Second is **quality**. It's crucial to ensure your wrapper provides access to a quality marketplace. Otherwise, you risk serving ads that conflict with your values, damage your brand, or even deliver malware to your users. To this end, any programmatic partner you work with in the header should have a mechanism for auditing the creatives they send you – ideally one that takes advantage of the latest technology but also gives humans final say on edge cases.

Your header bidding partners also need to live up to a **high performance** standard – both in terms of how they impact your page and how they participate in your auctions. Look at stats like mean and median bid response time, bid rate, win rate, and average CPM provided. You need partners who are frequently bidding on and winning impressions while also keeping your page load times reasonable.

Finally, your exchange partners should be able to bid against one another in **fair competition**. Publishers need a wrapper that will put every demand source on an even playing field. After all, why bother putting together a group of high-quality demand partners if only one or two of them ever has a chance to win the impression?

THE WRAPPER CODE OF CONDUCT

While we know that not every publisher will choose Prebid, we do believe there are a few key criteria a header bidding wrapper ought to fulfill -- even beyond connecting publishers with a large number of demand sources.

That's why we've established a code of conduct for header bidding wrappers. Our goal is to ensure that this technology will improve outcomes for everyone involved – not just for any one company, and not even just for publishers, but for advertisers and consumers as well.

For publishers to get the best results, they need to choose a wrapper that is:

1. Open and Transparent

No black-box auction dynamics, bid biasing, or hidden fees.

2. Easy

Publishers should only have to integrate once, and they should be able to customize the wrapper as they see fit.

3. Collaborative

Each wrapper should foster a community of contribution.

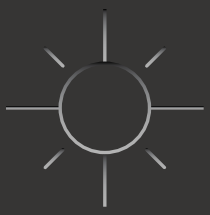

4. Agnostic

A publisher's wrapper needs to treat all adaptors and third party providers equally.

5. Performant

Wrappers cannot cause excessive latency, as this cripples user experience.

CHAPTER 2: SERVER-TO-SERVER HEADER BIDDING

Server-to-server takes the latency out of header bidding.

Wrapper technology made it easier for publishers to implement header bidding, but as they added more partners to the auction, some began seeing increases in latency.

In the traditional, client-side header bidding setup we described earlier, the publisher's web page makes ad calls from the user's browser. With each new demand source, the page needs to make additional ad calls, placing an even greater strain on the user's browser. Browsers have limited ports they can make calls with, so the whole process slows down if too many are trying to go through at once. Think of it as a giant traffic jam on a one-lane road. Generally speaking, a client-side auction can start to hurt page load times any time it includes seven or more demand partners. In essence, this means that the client-side setup was forcing publishers to choose between opening bidding to the optimal number of demand partners and protecting the user experience.

S2S may reduce cookie match rates and can be less transparent than client-side header bidding.

The technology that emerged to fix this problem is called server-to-server header bidding (S2S), a solution that reroutes the header bidding auction to reduce the strain on the browser. In S2S, the user's web page makes a single call to a third-party server, which then makes real-time calls to the publisher's various demand partners. Unlike browsers, servers aren't constrained by a small number of ports, so more calls can be made at once. To go back to that traffic jam analogy, implementing S2S is like building a six-lane freeway that can accommodate many more cars. Though it's frequently presented as a new technology, S2S is really just the reapplication of the protocol SSPs use to simultaneously call their DSP partners. In any event, S2S allows publishers to work with as many header bidding demand sources as they want, without harming the user experience.

Unfortunately, there are a few possible drawbacks to S2S that every publisher should be aware of.

The biggest problem with S2S is that it can make it harder for publishers to match their cookies with advertisers. Since media buyers rarely purchase programmatic ads without knowing who they're targeting at the other end, this issue can seriously damage publishers' ability to monetize.

The reason S2S lowers cookie match rates is that in client-side header bidding, the header makes ad calls from the browser, where cookies are stored, directly to programmatic demand partners. But in S2S, the auction takes place away from the browser inside a third-party server, and relies on user syncing between the third-party server and the various demand partners. As a result, it's tougher for advertisers and publishers to sync on user identity.

The other big issue with S2S is that it gives publishers less transparency than they get with client-side header bidding. Since client-side implementations can be executed via open-source code on the publisher page, media companies can easily monitor the auction for any signs of impropriety. By contrast, S2S auctions take place inside a third-party server that is essentially a black box for publishers, meaning they'd be unable to confirm your technology partner's revenue share fees or priotization logic of demand structures. However, open-source S2S solutions like the recently released Prebid Server don't have these issues, as publishers can simply look at the code themselves to confirm everything is running fairly.

This is why we recommend a hybrid solution that combines both S2S and client-side. With a setup like this, publishers can select which demand partners they want to call directly from their header bidding wrappers and which to call server-side. That way, publishers can maximize cookie matching with their client-side partners and increase bid density from server-side partners, while also managing latency and revenue optimization. This is possible today with Prebid's S2S solution, **Prebid Server**.

NOT SURE WHICH HEADER BIDDING SETUP IS RIGHT FOR YOU? HERE'S A BREAKDOWN OF THE KEY PROS AND CONS FOR EACH.

	CLIENT-SIDE HEADER BIDDING	S2S HEADER BIDDING
PRIMARY BENEFIT	Allows publishers to accept bids from multiple demand partners simultaneously	Allows publishers to accept bids from multiple demand partners simultaneously
EXECUTION	The webpage calls all demand partners directly	An adserver calls demand partners after receiving a single call from the webpage
LATENCY	Increases with every new auction participant	Remains stable as new partners are added
COOKIE MATCH RATES	100%	Match rates may fall
OPERATIONAL BURDEN	The publisher is responsible for managing the auction on its webpage	The vendor is responsible for managing ad calls and sourcing bids on its server
PUBLISHER OVERSIGHT	Depends on the provider. Open source solutions like Prebid allow publishers to comb through the code themselves for any issues, but most solutions are closed.	Most S2S solutions are even less transparent, as all the action takes place on a server owned by the provider. Open source solutions like Prebid Server, however, retain their transparency.

CHECK OUT THESE DIAGRAMS TO VISUALIZE THE DIFFERENCES BETWEEN THE CLIENT-SIDE AND S2S HEADER BIDDING PROCESSES.

CLIENT-SIDE HEADER BIDDING

When the user opens a web site, the site's header tags redirect them to SSPs **(1,2, and 3)**. The SSPs conduct an auction with the DSPs, and the DSPs respond with bids **(4 and 5)**. The SSP determines the winning bid value and returns it to the user's browser, which then pushes it to the publisher ad server **(5, 6, and 7)**. The ad server chooses a winning line item and serves the associated creative to the user **(8)**.

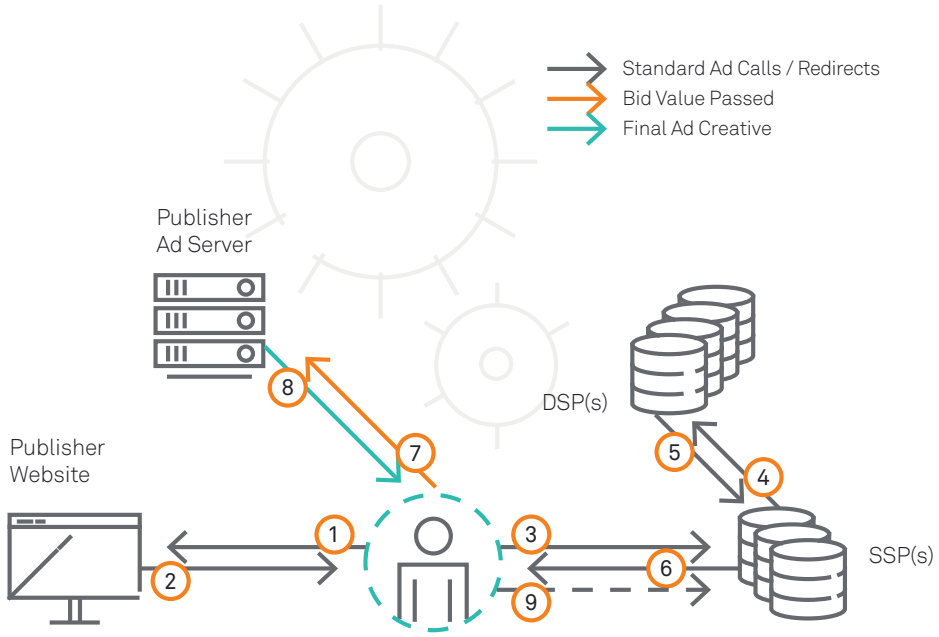

S2S/CLIENT-SIDE HYBRID HEADER BIDDING

When the user opens a web site, the site's header tags redirect them to the SSPs tagged on page **(1, 2, and 3)** just like in client-side header bidding. But then, the header sends ad calls via a third party server to SSPs and DSPs **(4a and 4b)**. Each of those demand partners send bids that go back to the tagged-on-page SSPs **(5, 6, and 7)**. Those SSPs send a winning bid value back to the user's browser, which then passes it into the ad server **(8 and 9)**. The ad server then decides on the winning line item and serves the creative to the user **(10)**.

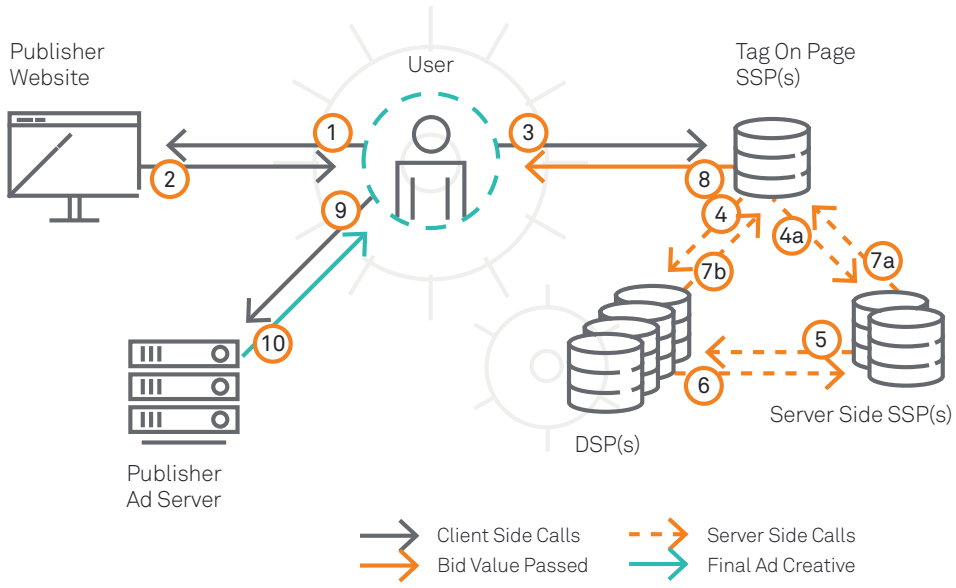

CHAPTER 3: NEW FORMATS

Header bidding is coming to video, mobile app, and native. Are you ready?

For the first two years of its lifespan, header bidding has been used primarily to sell display inventory. However, brands are increasingly shifting budgets toward emerging formats that deliver more relevant ad experiences. As publishers become more comfortable using header bidding, we expect that more and more of them will begin adopting it for video, mobile apps, and native.

Here, we'll break down how header bidding interacts with and enhances each of these fast-growing formats. For each format, we'll tell you why header bidding makes sense, how it can improve the format, and what you need to know to implement it successfully.

VIDEO

What to expect: Header bidding makes digital's hottest format faster, more transparent, and more lucrative.

Digital video is the buzziest format in advertising, and it's not hard to see why. With each passing year, consumers spend more time watching video, and advertisers are following the eyeballs.

According to eMarketer, today's average adult spends 72 minutes a day watching online video, up from just 46 minutes in 2013. Meanwhile, video ad spend is expected to reach $17.56 billion by 2019 -- a staggering increase from the $7.68 billion marketers spent in 2015. And if anything, advertisers want to be spending more. In a recent study analyzing price differences between different formats, eMarketer found that on one SSP, some video ads were reeling in as much as $11.50 per impression, compared to just $1.10 for display.

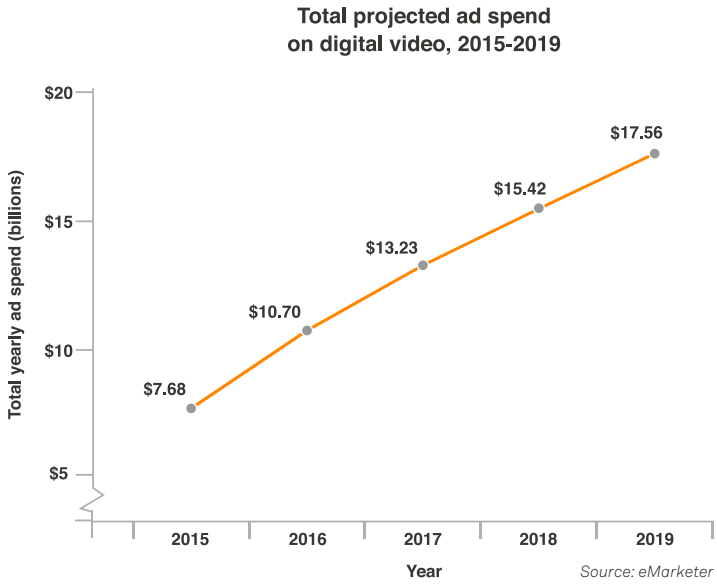

If you think that's impressive, just imagine the kind of money publishers will be able to make once they add header bidding to the mix. For instance, if video publishers start getting the 50% CPM boost that some media companies saw in the programmatic display market after adopting header bidding, those $11.50 CPMs will grow to more than $17. In fact, we're already seeing this sort of performance with our clients. When our client Playwire – an HTML5 video player and monetization platform – implemented header bidding, its CPMs rose as high as $30.

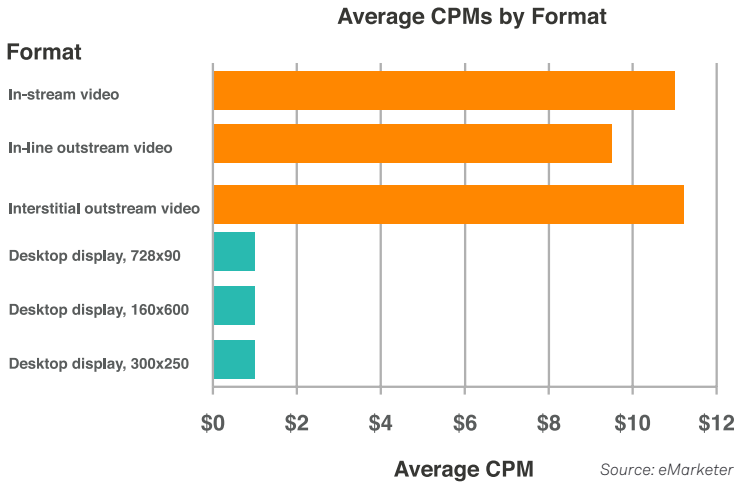

HEADER BIDDING'S POTENTIAL IMPACT IN THE VIDEO SPACE CAN BE BOILED DOWN TO THREE MAJOR BENEFITS:

Superior monetization

Compared to display, a bigger share of video inventory is sold through direct deals, mainly because it's in such high demand – especially for premium publishers. Unfortunately, the share of video that is sold programmatically is usually distributed using the same waterfall process that previously cost publishers in the display market.

By opening auctions to every demand partner a publisher works with, header bidding increases competition and drives video CPMs even higher. The added competition also gives publishers a better sense of what their video inventory is worth on the open market, allowing them to negotiate better direct-sold deals. This is important for all your inventory, but especially for premium inventory like video.

Lower latency

Video publishers are highly sensitive to page latency, as viewers are likely to exit the page if their video takes too long to load. **In a study of over 23 million video playbacks**, researchers found that viewers start to leave after just two seconds of video load time, with each additional one second of latency causing another 5.8% increase in user abandonment rate. Though load times have gotten much better in recent years, the user experience is frequently undermined by the publisher waterfall setup. In the waterfall, every new demand partner the adserver calls can add one to three seconds of load time. Even worse, some demand partners create additional latency by holding their own secondary auctions when they can't source an ad for an impression.

By contrast, header bidding replaces the waterfall with a single, simultaneous auction across all of a programmatic demand partners. This setup also gives publishers the opportunity to set global timeouts across their header bidding partners, ensuring that the auction is capped at a time that doesn't hurt user experience.

Transparency and Market Data

The waterfall system prevents publishers from finding out how their demand partners really value their inventory. Since partners lower down the chain frequently miss out on the opportunity to submit bids, it's impossible to know what they would have been willing to pay for impressions. And when technology providers don't offer insight into auction mechanics, publishers can only hope that their partners aren't biasing certain sources.

By opting for an open-source wrapper solution like Prebid, publishers can monitor the header bidding auction logic for themselves. Since Prebid returns a list of every bid submitted for each impression, you'll be absolutely positive that you're getting the highest possible yield.

How it works: For all that video header bidding brings to the table, it's fairly simple to implement.

From a technical perspective, video header bidding isn't all that much different from display. While there isn't a header inside video players, it's important to remember that the video player is just another piece of Javascript on the page. Once you **integrate your video player** with your header bidding wrapper, the header can easily pass winning bids into the player.

From there, implementation becomes a simple, three-step process. First, you create the video ad units on your page. Then, you add those units to your wrapper to source demand. The last step is to add new line items for video inside your adserver -- just like you would in any other ad format.

Best practices: Three tips for reducing latency and improving the ad experience.

1. Set a timeout.

Keep load times to a minimum by setting a timeout inside your header bidding wrapper. This will act as a deadline for all of your demand partners to return a bid -- we recommend 700 milliseconds.

2. Move the auction away from the play button.

Unlike in display advertising, most web pages don't start calling for video ads until the user hits the play button on the video player. This creates latency, as the user cannot watch their content until the page has sourced the ad.

Minimize this delay by setting up your page to auction off video inventory as soon as the page loads, rather than waiting for the user to click the play button. Then, your page can pre-cache the XML file of the winning video ad on the user's device. This makes it so the ad will be retrieved extremely quickly once the user actually reaches the ad slot.

3. Manage your video playlists.

A word of caution: If you have a playlist and you make an ad request from the page level, your ads will be aligned with the metadata and the content of the page, but not necessarily with every video in the playlist.

Instead, be sure to use the metadata associated with each video in the playlist. This way, the ads will align with the content in each clip.

MOBILE APP

What to expect: Mobile header bidding provides fresh demand, stronger monetization, and a much better user experience.

Header bidding has had a huge impact on display advertising, with some publishers increasing revenue by 50% or more. With that in mind, it's no wonder that the majority of the web's top publishers have adopted this powerful new monetization technique. So, what's the next stop for header bidding? We think it could be the mobile app.

The numbers don't lie: Apps are becoming the best place to put mobile ads in front of users. According to **recent research**, consumers are spending more of their time with mobile apps each year, going from 1 hour, 55 minutes a day in 2015 to 2 hours, 25 minutes in 2017. At the same time, mobile web usage has remained relatively stagnant, suggesting that apps are driving the overall increase in mobile internet usage. Meanwhile, mobile adoption **continues to grow**, particularly in the developing world, where apps are also extremely popular.

Projected average time per day spent with mobile app and mobile web, U.S. adults, 2015-2019

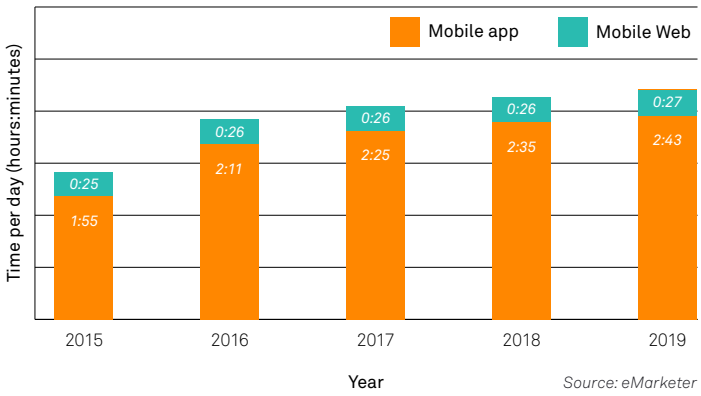

Source: eMarketer

Of course, in-app advertising is also doing pretty well in the here and now. According to **eMarketer's latest report**, mobile app CPMs are only slightly lower than desktop CPMs, and even this discrepancy is mostly the result of brands undervaluing in-app advertising. For advertisers, in-app provides superior targeting across the board, while standout formats like in-app interstitial deliver near-guaranteed viewability.

For publishers, header bidding – or technically speaking, the mobile equivalent, given that apps don't have headers – is about to take this already-powerful medium to the next level. Here's what you can get from implementing it:

Incremental revenues and higher fill rates

Just like on desktop, header bidding increases monetization and fill. By expanding the number of partners who can bid on your app's ad units, header bidding can increase incremental yield by up to 15%. In the long run, that heightened competition gives you a better sense of your inventory's true market value and helps you get fairer prices on every impression.

A greater variety of supply and demand

Mobile apps encompass a wide range of publishers that serve users all over world. Just think, a mobile app can be a news provider like The New York Times, a game like Clash of Clans, or a dating app like Bumble or Grindr. But despite this incredible variety of both users and contexts to reach them, most apps are dominated by app install ads. That means there are untapped opportunities to personalize mobile app ads and provide a more relevant experience to every user.

With the right header bidding partner, you can build a much more diverse demand pool by accessing big brands and global advertisers. With their enhanced targeting capabilities, your users will get ads that better match their preferences, as well as the unique content of your app. These advertisers can also help you better monetize your international audience, which is important considering the global reach of mobile apps.

A better user experience

In a world where apps immediately **lose 80% of the users** who download them, it's hard to understate the importance of a quality user experience. But when app developers try to sell their inventory programmatically today, users too often have to wait through long load times as their impressions are passed to different ad networks during the mediation process.

By offering your inventory directly to exchanges, mobile header bidding heavily reduces reliance on mediation and creates a much faster mobile app. It's not just about monetization -- header bidding could provide the retention boost that makes your app successful.

How it works: Mobile header bidding pre-caches ads in advance, ensuring that they'll be ready to go when you need them

Similar to the video format, mobile apps do not have a header. As such, implementation is a bit different from what you might be used to in desktop display.

In desktop header bidding, the Javascript in the header delays calls to the ad server for a bit while it calls the publisher's demand partners. But in mobile apps, the script works in the background the whole time. While the user browses the app, the script anticipates possible ad opportunities and pre-caches the winning bids and creatives in advance. Then, when the user reaches the ad, the creative is served with little to no delay. Latency is the biggest detriment to user experience on any format – not just mobile apps – so the ability to pre-cache winning ads is a huge benefit of mobile header bidding.

Best practices: Three tips for speeding up implementation, managing the user experience and accessing quality demand.

1. Use S2S if you want to work with multiple demand partners.

Server-to-server header bidding doesn't just create a better user experience for consumers while they're in your app; it also makes their lives easier before they open it.

In a client-side implementation, app publishers have to force their users to update the app every time they add a new demand partner. A server-side solution like **Prebid Mobile**, on the other hand, lets you add new demand partners without bothering your users.

2. Think hard about your demand partners.

With header bidding, you no longer have to worry as much about mediation, which means you can shift focus to choosing the best possible demand sources for your inventory. Look for partners who give you access to advertisers who match both your audience and app content.

3. Educate your developers.

Since most app developers don't have prior ad ops experience, implementing header bidding can be a tough assignment. Make sure they know the basics about how ads are requested and rendered on page so that they can design a solution that doesn't add too much weight to your app. No mobile app has to make due with a waterfall setup anymore -- it's just a matter of helping developers learn what the best solution is and how to implement it.

NATIVE

What to expect: Header bidding could be the match that lights the fuse for programmatic native.

At a time when both advertisers and publishers are increasingly improving user experience, native advertising is poised for exponential growth. These unique ad units advance that goal in a few different ways:

Native ads blend in with the look and feel of a given web site, providing maximum relevance to the consumer without disrupting what they're doing online.

Native ads generally have more content – and not just that, but more engaging content as well, further increasing the relevance to users.

Native ads are lightweight and rely on very little Javascript, which means they load faster. They're also responsive, meaning they look great and load smoothly no matter what device the user is browsing on.

This creates a win-win situation for all involved: consumers get a better user experience, publishers enjoy a more engaged audience with higher CPMs and advertisers deliver more effective ads that match the content people are consuming.

Just this year, total native ad spend is **expected to grow 36%** to $22.1 billion. However, the vast majority of this spend is coming through pre-negotiated direct deals, meaning almost everyone is missing out on programmatic's superior targeting and scale, as well as the increased competition it brings to the table. As a result, programmatic native CPMs are still relatively low. While there are a few walled gardens that do programmatic native well, most providers function a lot like ad networks did ten years ago, taking huge cuts without adding much value. Certainly, programmatic native requires more technical work than traditional display, as advertisers and publishers must work together to ensure that the creative fits the unique style of each site it runs on. Still, this challenge is far from insurmountable. In fact, we believe that the programmatic native market will take off once publishers begin adopting header bidding.

Here's why: By adding more players to the auction, header bidding can open more points of competition in the native programmatic market. As a consequence, existing programmatic providers will be compelled to lower their take rates, making programmatic native a more attractive option for publishers.

And once that happens, the entire digital ecosystem will be better off. As header bidding creates a real-time marketplace for native ads, more and more brand dollars will flow into the space, creating a superior user experience filled with better, less interruptive advertising. Not to mention, this is premium inventory that's already fetching high prices in direct deals. Bringing more buyers into the space will increase competition and drive that value even higher for publishers.

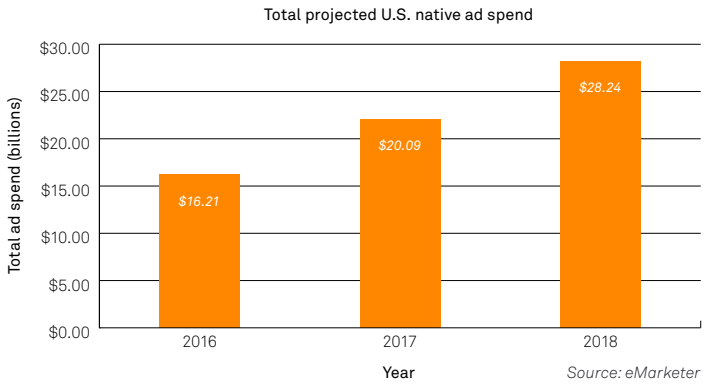

How it works: Native header bidding takes a little more work, but it gives publishers greater control over the user experience.

Right now, many publishers favor direct deals in native because they think of the format as highly customized. But in reality, native ads can be delivered programmatically via header bidding the same way display ads are.

The biggest difference is that in native, publishers receive components of the ad separately -- the creative, the copy, etc. -- and then assemble it on the page themselves. This process takes a little more time and effort, but it also gives publishers more control over how the ad renders, and thus, more control over user experience as well. In addition, since native ads are built to fit into a page's existing infrastructure, they use less Javascript and have fewer tracking mechanisms than other formats. Those two factors mean that native ads present a lower risk of data leakage in header bidding.

Best practices: In order to maximize native's value, publishers must adopt a programmatic mindset.

A stronger programmatic marketplace can kickstart native advertising, and header bidding is the key. But in order for that to happen, publishers need to change the way they think about this innovative format. Rather than conceiving of the format solely as a place for small, bespoke campaigns, publishers must start giving thought to how they can use native to also deliver scale.

To be clear, there will always be room for tailored, direct campaigns that perfectly mesh with a publisher's brand. But in order for us to realize native's full potential, we need to build a marketplace where more generic, low-fidelity native ads can be bought and sold in large quantities. Fortunately, native uses lightweight units that can be broken down into their component parts and repackaged by the publisher to match the user experience. As such, the marketplace we need is entirely within reach. All it will it take is for publishers to start thinking of the format a little bit differently.


WHAT'S NEXT FOR
HEADER BIDDING?

THIS IS ONLY THE BEGINNING

Ultimately, our goal today is the same as it was when we first began developing our header bidding products. In helping publishers open their inventory to the full scope of the market's demand, we strive to improve yields by returning CPMs that accurately reflect what each impression is really worth. While header bidding technology will continue to evolve over the next several years, this will never change.

If anything, our commitment to this objective will only grow stronger as header bidding becomes a more effective monetization solution. In S2S header bidding, publishers are adopting a powerful new tool to make their auctions fairer, more accessible, and better for the user experience. Meanwhile, header bidding's continued expansion into video, mobile app, and native is empowering publishers to maximize yields across an even greater portion of their inventory and give their users a better experience across all the formats they operate in.

We're also thrilled by the positive impact that header bidding has had on ad server technology. This has played out in a few ways. By enabling publishers to hold a unified auction with all of their demand partners, header bidding has flattened the waterfall and forced the ad server to compete its demand against that from other sources, much to publishers' excitement. In addition to making auctions more open, header bidding is also changing the role of the ad server. For instance, in the last year, Prebid switched from running an in-wrapper auction to simply passing all bids into the ad server. So rather than just placing ads in the right place, our ad server is becoming the central yield management hub and site of the final auction for all impressions, imbued with powerful decisioning technology. Moving forward, any ad server that doesn't embrace header bidding and allow for more open, transparent auctions will be at a distinct disadvantage.

Taking a higher level view, header bidding is creating a better internet for everyone. As technology improves and adoption rates increase, we expect header bidding to become a growing part of a more robust digital marketplace -- one where publishers can get fair, market-determined prices for their impressions, and where advertisers can deliver relevant ads to the right audiences. By providing superior results for advertisers, this marketplace will encourage greater investment from brands, giving digital publishers more money to produce the quality content consumers love.

It's an exciting future, and one we are very much looking forward to building alongside our publisher partners. Our hope is that this whitepaper has answered some of the questions you've been asking, and that this information will help you chart your path forward as the ad tech ecosystem continues to grow and change. If there's anything else you'd like to know, please do reach out to an AppNexus representative. We'd love partner up and help you get even better results from your header bidding monetization stack.

Want to learn more about how header bidding can help your business? Contact us here