avenga

The AdTech Book

Field-Tested Knowledge and Practices from the AdTech Frontlines

New 2026 Edition

About Avenga

Avenga is a global engineering and consulting partner, trusted by leading AdTech, MarTech, and media companies to design, build, and scale high-performance platforms.

From custom demand and supply-side solutions to data management, consent, and attribution systems, we bring deep domain expertise and cross-functional teams to address complex technical challenges across the advertising ecosystem.

Throughout this book, you'll find real-world examples of our work — each one showing how we help clients transform vision into product, faster and smarter.
If you're facing a build, scale, or integration challenge in AdTech, we'd love to hear from you.

Contact us at avenga.com
to start the conversation.

1. Basics

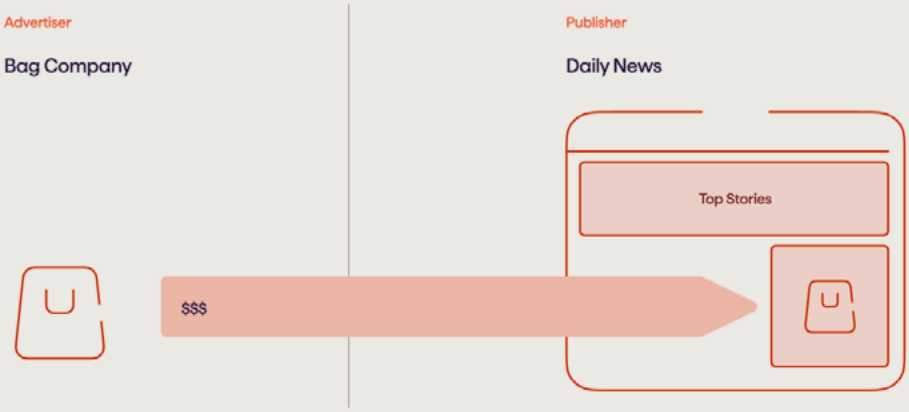

Before diving into the complexities of the AdTech ecosystem, it's essential to understand its core components.

This chapter covers the foundational terms, tools, and concepts that power digital advertising — from ad slots and impressions to pricing models like CPM, CPC, and CPA.

Whether you're new to the space or need a quick refresher, this is where your AdTech journey begins.

AdTech

AdTech, short for *advertising technology*, refers to the software, tools, and processes used to create, run, manage, measure, and optimize digital advertising campaigns. Examples include:

- Demand-side platform (DSP)
- Supply-side platform (SSP)
- Advertising networks
- Creative optimization processes
- Digital-out-of-home advertising
- Mobile app advertising
- Budget capping
- Media-buying processes, etc.

A well-built and optimized AdTech stack helps advertisers reach their desired audience at the right time, in the right place, and on the right device — for example, during business hours at home on a laptop or late at night in a city center on a mobile phone.

Publishers, on the other hand, gain tools to monetize their assets — primarily ad space. Here's a list of just some of the things you can achieve with an AdTech stack tailored to your needs:

- Decide which ads to display to specific user groups based on the advertiser's targeting criteria
- Deliver ads across multiple channels (web and in-app mobile) and devices
- Optimize campaign performance for advertisers and yield for publishers
- Collect user data and build audiences
- Generate measurement and analytics reports
- Manage billing and media-buying processes

The online advertising industry has its own terminology and an endless supply of acronyms. Here are some of the most important ones.

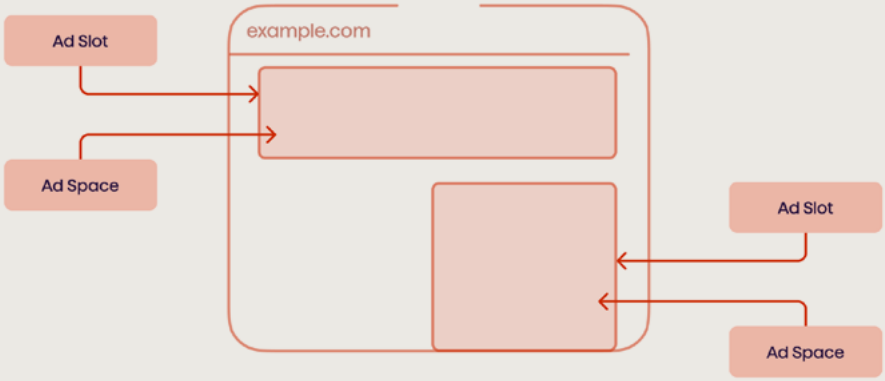

Ad Slot

An ad slot is the specific space on a website reserved for displaying ads. Within each ad slot, an ad tag communicates with the ad server to load and render the actual advertisement.

Ad Space

Ad space is the actual impression available in the ad slot. While the two terms are often used interchangeably, the key distinction is that the ad slot is the container, whereas the ad space is the area inside it that displays the ad.
To help clarify this point, think about a billboard. The actual billboard would be the ad slot and the white section inside the billboard where the ads are placed would be the ad space.

Inventory

Inventory, aka ad inventory, is the name given to all the ad space available on a website. Although the terms inventory and ad space are sometimes used interchangeably, inventory usually describes the overall supply of available space. There are three main types of ad inventory: premium, remnant, and long-tail.

- **Premium inventory** is the publisher's most valuable space. It is typically found on well-known publisher sites and high-traffic pages — such as homepages — or in highly visible placements, like the top of a page.
- **Remnant inventory** is inventory that a publisher has been unable to sell directly to advertisers through direct deals and other primary inventory-monetization channels. It basically is the leftover inventory publishers are willing to sell for less than their standard price.
- **Long-tail inventory** is found on small sites and blogs. A large chunk of this type of inventory is sold via Google AdSense. Long-tail sites sometimes seek better monetization alternatives to AdSense, for example, by joining affiliate programs and networks.

Creative

A creative is the file that contains the actual advertisement a user sees or hears. The most common types of creatives are:

- Graphical files (either static or animated)
- Video files
- Audio files

Creatives can come in various formats, including GIF, JPEG, and HTML5 (previously Flash) for graphics, and MOV, FLV, and MP4 for video.

Click

A click is counted when someone clicks on an ad, even if the person doesn't reach the advertiser's website, for example, if it's temporarily unavailable.

Impression

An impression — sometimes called an ad view — is counted each time a creative is served. Importantly, if a user refreshes the page and sees the same ad again, a new impression is still recorded.

Viewable Impression

A viewable impression is a metric used to determine whether an impression was actually seen by a real human or whether it was "seen" by a bot or hidden from the user's view. For example, at the bottom of the page where the user doesn't scroll.

Conversion

A conversion is recorded whenever a user completes a specific goal defined by an advertiser or marketer.

Examples of conversion goals include purchasing a product, signing up for an online service, downloading a resource such as an ebook, or filling out a contact form on a landing page.

Landing Page

A landing page is a web page where an online user "lands on" after clicking on an ad or link.

Unlike regular pages on an advertiser's website, landing pages are built with a single objective: driving conversions. Their purpose is to get users to complete a specific goal, such as making a purchase, signing up, or filling out a form.

Advertising campaigns often send traffic to landing pages to maximize conversions, particularly in direct-response or prospecting campaigns.

CPM vs CPC vs CPA

Cost Per Mille (CPM) — The CPM pricing model refers to the cost per 1,000 impressions; the word mille is Latin for one thousand. The reason we use thousands instead of individual impressions is that the price per impression is very small compared to clicks or conversions, and it would be inconvenient from an accounting perspective if the advertiser paid $0.002 per impression.

Cost Per Click (CPC) — Cost per click is a pricing model used to express how much each click on an ad or link would cost the advertiser. If an advertiser buys ad inventory from a publisher on a CPC basis of $1.10, then every time a visitor clicks on the ad, the advertiser is charged $1.10.

Cost Per Action/Acquisition (CPA) — Under this model, the publisher or affiliate is paid only when a user converts (e.g., purchases a product or submits a lead form) after viewing or clicking on an ad. CPA is commonly used in affiliate networks but is less widespread than CPM or CPC.

2. Technology

This chapter explores the core technologies that power modern digital advertising. From DSPs and SSPs to DMPs, ad servers, and beyond, you'll get a clear understanding of the platforms, players, and processes that automate, optimize, and scale programmatic campaigns in real time.

Demand-Side Platform (DSP)

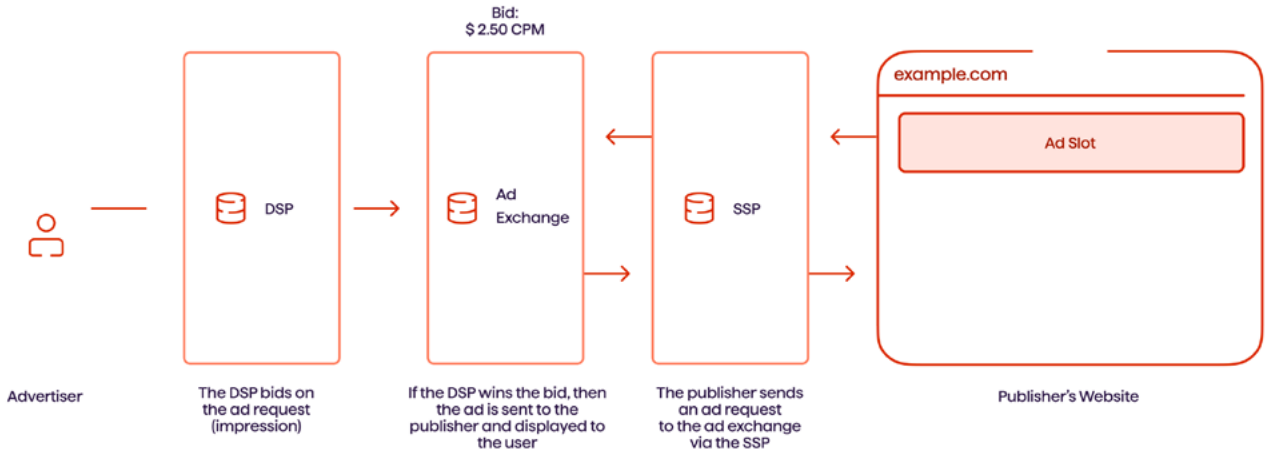

What Is a Demand-Side Platform (DSP)?

A demand-side platform (DSP) is a platform that allows media buyers (advertisers and agencies) to run advertising campaigns and buy inventory from various ad exchanges and supply-side platforms (SSPs) through a single interface.

DSPs utilize real-time bidding (RTB) technology to facilitate the efficient and targeted buying of ad space across various platforms, including websites, mobile apps, and social media.

To help improve targeting and enhance media buys, DSPs often utilize data from data-management platforms (DMPs) and data brokers.

What Are the Main Functions of a DSP?

- Create, run, and manage multiple campaigns simultaneously across SSPs and ad exchanges, all from a single user interface
- Auto-optimize campaigns using algorithms and machine learning to improve ROI, viewability, conversions, cost per action (CPA), and more
- Leverage data from customer data platforms (CDPs) and data management platforms (DMPs) to enhance targeting
- Provide real-time reporting through advanced analytics

Case Study

Proprietary DSP with Reinforcement Learning — We built a custom DSP with a Softmax Multi-Armed Bandit algorithm to optimize creatives based on conversions.
The client achieved significantly higher campaign performance and more new game installations.

We used 17 technologies to create this project.

How Does a DSP Work?

First, an advertiser sets up a campaign in a DSP, defining targeting criteria and uploading creatives. When an impression becomes available, the DSP receives a bid request from the SSP or an ad exchange containing contextual and user-related information. The DSP's bidder then decides whether to bid, how much to bid, and which creative to serve. If the bid is successful, the creative is delivered through the DSP's ad server.

This process occurs every time a user visits or refreshes a page, with the real-time bidding (RTB) cycle typically taking about 100 milliseconds.

What Are the Main Components of a DSP?

- Bidder
- Ad server
- Campaign tracker
- Banker
- Reporting database
- User profile database
- User interface
- API connectors
- Creative management tools

Key Integrations:

- SSPs and ad exchanges
- CDPs, DMPs and data brokers
- Ad verification and creative optimization platforms
- Meta-DSP

Supply-Side Platform (SSP)

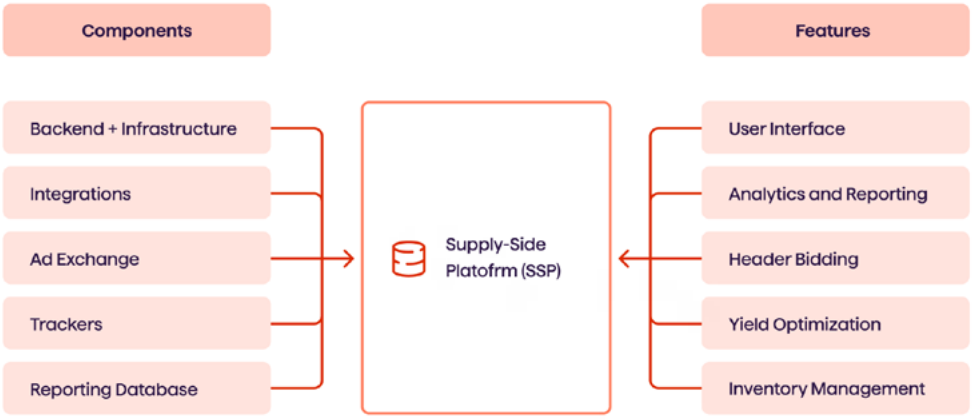

What Is a Supply-Side Platform (SSP)?

A supply-side platform (SSP) is an advertising technology platform that allows publishers — such as website owners, mobile app developers, or other digital property holders — to manage, sell, and optimize their ad inventory. Through an SSP, publishers can connect to real-time bidding (RTB) marketplaces and programmatic deals, selling display, video, and native ad space on an impression-by-impression basis.

SSPs streamline ad operations, automate sales, and help maximize revenue by making inventory simultaneously available to multiple demand sources, including DSPs, ad exchanges, and direct buyers.

Case Study

Premium Supply Platform (PSP)
— We developed a deal-based SSP supporting multiple programmatic deal types. The client gained a flexible, premium platform to maximize monetization across open and private deals.

We integrated this solution with 2 big providers: EngageBDR and RTBiQ.

What Are the Main Functions of an SSP?

- Manage and organize ad inventory, including formats, sizes, and targeting options
- Connect to multiple ad exchanges and DSPs to maximize competition for each impression
- Run real-time auctions and select the highest bid in milliseconds
- Provide ad quality controls, fraud detection, and brand safety protections
- Support audience targeting, frequency capping, and privacy-compliant data use
- Deliver real-time and custom reporting for revenue and performance tracking
- Support both client-side and server-side header bidding to increase competition and yield
- Allow operational flexibility through macros, overrides, and custom settings
- Optimize yield with dynamic floor pricing, packaging strategies, and data-driven analysis

How Does an SSP Work?

A publisher defines its available ad space and sets parameters such as format, size, targeting criteria, and floor prices. When a user visits a website or app, an ad request is sent through the SSP to connected ad exchanges, DSPs, and demand partners. The SSP conducts a real-time auction, evaluates bids based on price, relevance, ad quality, and publisher rules, and determines the winner.

The winning bid is returned to the publisher's site or app, and the winning creative is served to the user. This process happens under a second for every eligible impression.

What Are the Main Components of an SSP?

- Auction engine and bidder infrastructure (often cloud-hosted)
- Integrations with ad servers, DSPs, ad exchanges, and data platforms
- Ad exchange functionality for direct inventory trading
- Header bidding module for simultaneous demand competition
- Data collection systems for audience and performance insights
- Reporting database for real-time and historical analytics
- User interface for campaign setup, monitoring, and yield management

Key Integrations:

- Ad exchanges, DSPs, and direct buyers
- Data platforms (DMPs, CDPs, or clean rooms)
- Ad servers
- Yield optimization tools and header bidding frameworks (e.g., Prebid)

Data-Management Platform (DMP)

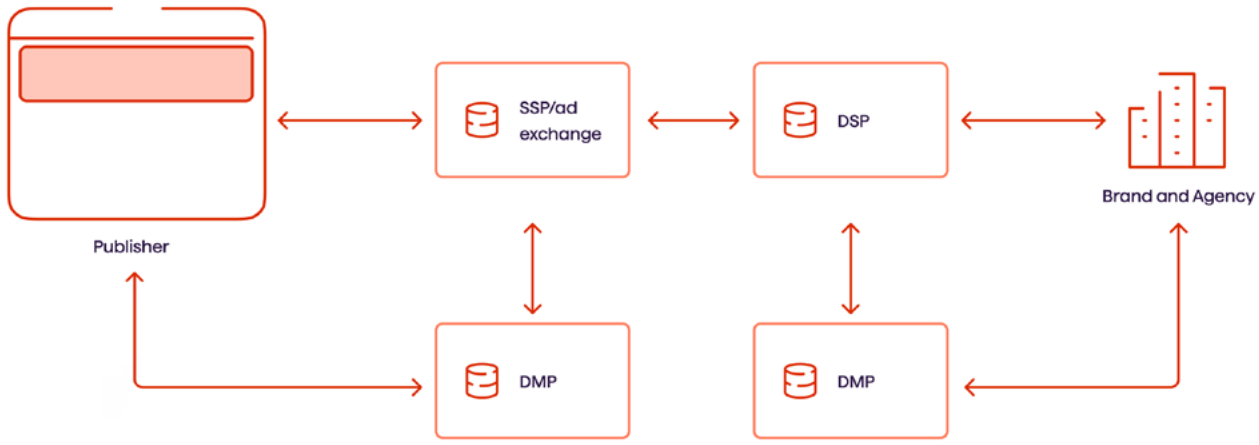

What Is a Data Management Platform (DMP)?

A data management platform (DMP) is software that collects, stores, and organizes data from multiple sources, including websites, mobile apps, and advertising campaigns. Advertisers, agencies, and publishers use DMPs to enhance ad targeting, conduct advanced analytics, perform look-alike modeling, and support audience extension.

What Are the Main Functions of a DMP?

- Data collection
- Data organization
- Audience building and segmentation
- Activation
- Targeted advertising and personalization

How Does a DMP Work?

A DMP begins by collecting data through integrations — either server-to-server or via an API — with other AdTech and MarTech platforms such as DSPs, ad exchanges, SSPs, and customer relationship management (CRM) systems.
It can also collect data by placing a tag (a JavaScript snippet or HTML pixel) on an advertiser's or publisher's website.
Once collected, the data goes through a series of processes, including:

- Data normalization and enrichment
- Profile merging
- Profile building
- Audience creation
- Data storage
- Data activation
- Data segmentation (data classification) and taxonomies
- Analytics and reporting

What Are the Main Use Cases of a DMP?

- Improving targeting for online media campaigns
- Audience extension
- Onsite personalization (this capability is often shifted to CDPs or personalization engines, because they work with real-time, first-party data better than legacy DMPs)
- Content and product recommendations

Ad Server (First- and Third-Party)

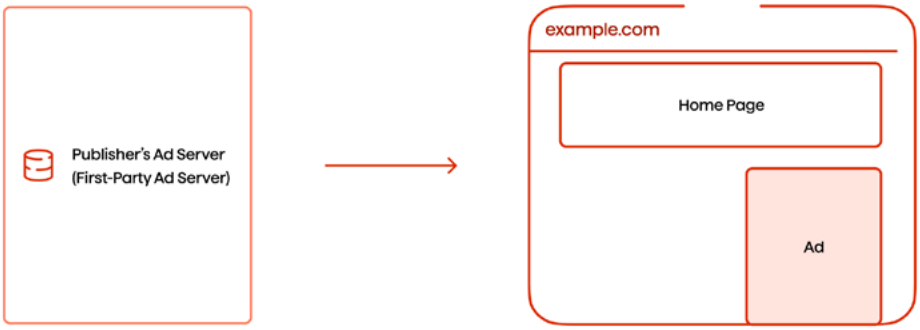

Key Integrations:

- Advertising platforms: DSPs, SSPs
- CRM systems
- Marketing automation tools
- Analytics platforms
- Data warehouses
- Social media platforms
- Integrations with publishers, clean rooms, and consent management platforms (CMPs) are becoming standard (CMPs) are becoming standard What Is an Ad Server?

What Is an Ad Server?

An ad server is an AdTech platform that decides which ads to display, delivers them to users, and collects and reports data on impressions, clicks, and other performance metrics. Much like WordPress is used to manage a website's content, an ad server is used to manage and deliver advertising on a site or app.

Publishers, advertisers, and agencies rely on ad servers to run multiple campaigns, often connecting them with other AdTech platforms such as DSPs and SSPs.

What Are the Main Functions of an Ad Server?

- **First-party ad server (publisher's ad server)** — fills the ad slots on a website by matching ads from direct campaigns, real-time bidding (RTB) auctions, and other media-buying processes. It also forecasts future inventory availability and predicts campaign performance using current and historical data.
- **Third-party ad server (advertiser's ad server)** — tracks the performance of the whole campaign across all publishers in a single system. It measures campaign reach (accounting for co-viewership across publishers) and verifies the accuracy of publisher-reported data.

How Does an Ad Server Work?

1. The ad request. When a user visits a web page, the browser loads the ad tag and sends a request to the ad server. This tag is a snippet of JavaScript containing details like the ad placement's name and position, so the server knows which ads are relevant to return.
2. Compiling campaigns. The ad server compiles a list of all campaigns assigned to the requested placement, taking into account each campaign's scheduled flight dates.
3. Checking eligibility. The server then filters out ineligible campaigns - such as those targeted only to mobile devices or specific geographic regions.
4. Selecting the ad. From the pool of eligible campaigns, the server chooses one ad to serve. The selection can be random, weight-based, or influenced by factors like the advertiser's bid.
5. Delivering the ad code. The server responds with the HTML/JavaScript code needed to render the ad. The browser executes this code to display the ad on the page.

6. Tracking impressions and clicks. Once the ad appears, the server logs an impression. If the user clicks the ad, a call is made to the server, allowing it to record the click and calculate performance metrics like click-through rate (CTR). Click/impression tracking often includes also viewability measurement, invalid traffic detection, and conversion tracking, often integrated with advertiser analytics.
7. Monitoring performance. The ad server maintains a database of all tracked events and provides performance reports through its user interface. These reports help advertisers and publishers evaluate results and serve as the basis for billing.

What Are the Main Components of an Ad Server?

- Ad decision engine
- Ad planning and scheduling
- Creative management and storage
- Ad delivery/service module
- Targeting and personalization
- Tracking and reporting
- Inventory management and forecasting
- User interface

Key Integrations:

- DSPs
- SSPs
- Ad exchange
- DMPs
- Creative optimization platforms (DCO) for personalization
- Verification vendors
- CDPs for first-party audience targeting

Ad Network

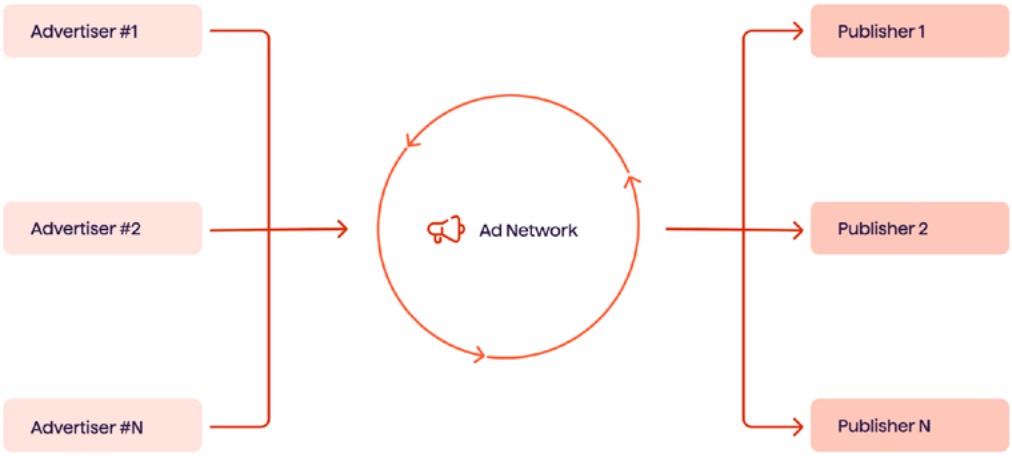

What Is an Ad Network?

An ad network is a technology platform that brokers deals between groups of publishers and advertisers. It aggregates unsold publisher inventory and offers advertisers a consolidated, typically lower-cost pool of impressions on a CPM basis.

Although ad networks are primarily used to help publishers sell remnant inventory, they also provide several advantages for advertisers:

- **Scale** — Advertisers can purchase inventory from multiple publishers through a single ad network and centralize campaign reporting.
- **Time savings** — Campaigns can be set up once, eliminating the need to negotiate or sign insertion orders with individual publishers.

- **Campaign reach and measurement** — Ad networks track reach across publishers and apply frequency capping at the campaign level.
- **Monetization** — Publishers can monetize unsold inventory that wasn't secured through direct deals.

Many ad networks specialize in certain types of inventories, e.g.:

- Premium ad networks — Offer inventory from top-tier publishers (e.g., *The New York Times*).
- Vertical ad networks — Focus on specific content categories such as business, technology, automotive, or fashion.
- Specialized ad networks — Concentrate on a single channel type (e.g., mobile, video, or native).
- Performance and affiliate ad networks — Operate on revenue-share, CPC, or CPA pricing models.

The targeting and decision-making processes in ad networks work in a similar way to those found in ad servers. By setting targeting criteria for a campaign, an advertiser can choose which web traffic is relevant for them. Common targeting criteria in ad networks include:

- Run on network (RON) — run on all sites in the ad network
- Run on site (ROS) — target specific domains/ publishers in the ad network
- IAB categories
- Geolocation
- Keywords (context)
- Time of day
- Browser type / OS
- ... and many others

What Are the Main Functions of an Ad Network?

Ad networks aggregate unsold publisher inventory and provide advertisers with a consolidated, typically lower-cost pool of impressions sold on a CPM basis.

How Does an Ad Network Work?

- An ad network aggregates inventory from many publishers to provide advertisers with sufficient scale, typically on an auction basis.
- Advertisers can set up campaigns directly through the ad network's campaign management panel, or they may implement tracking pixels from a third-party ad server for verification and consolidated reporting — especially when running campaigns across multiple ad networks or in direct deals with publishers.
- In this process, the advertiser defines campaign parameters such as targeting, budget, and frequency caps, while the publisher installs the ad network's tags on their site, either by embedding them directly into the page or via a first-party ad server.
- Once the ad is live, advertisers can rotate multiple banners on the publisher's site through the ad network's campaign management panel without needing to contact the publisher.

Key Integrations:

- Publisher websites and apps via JavaScript tags, header bidding adapters, iFrames, CMS-specific plugins, and mobile SDKs
- Ad servers
- SSPs
- DSPs
- Header bidding adapters
- Analytics and attribution providers
- DMPs, CDPs, and marketing platforms through APIs or direct integrations
- RESTful APIs

Ad Exchange

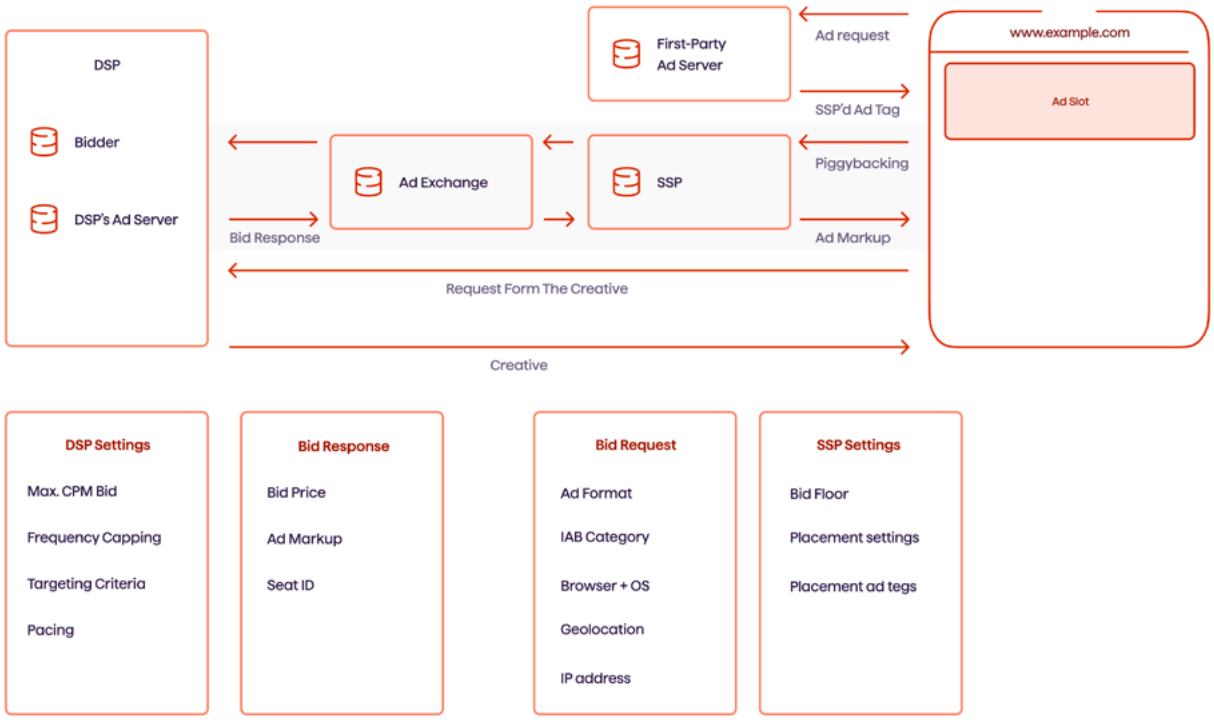

What Is an Ad Exchange?

An ad exchange is a dynamic technological platform that facilitates the buying and selling process of available impressions between advertisers, who place their bids via DSPs, and publishers, who sell their inventory via SSPs or directly with the ad exchange.

How Does an Ad Exchange Work?

User visit — When a user opens a website or app, the publisher's SSP sends a request to the ad exchange for an available impression.
Real-time bidding (RTB) — The ad exchange initiates a real-time auction in which advertisers, via their DSPs, can bid on that specific impression.
Bidding and targeting — Advertisers leverage audience data (e.g., demographics, location, browsing history) to set their bid price and target the most relevant users.
Winning bidder — The ad exchange selects the highest bidder, and the corresponding ad is chosen for delivery.
Ad served — The publisher's ad server delivers the winning ad to the user's device.

Key Integrations:

SSPs and DSPs
Analytics platforms and optimization tools
Anti-fraud systems

The Walled Gardens

Google amazon Meta

The AdTech ecosystem can be divided into two groups: independent AdTech companies and walled gardens. A walled garden means a closed ecosystem where a platform or technology provider controls access to its users, data, and advertising inventory.

In AdTech, you'll often encounter the term GAMA — Google, Apple, Meta, and Amazon — considered among the most prominent walled gardens. These platforms keep their audiences and data within their own ecosystems, requiring brands to use their advertising tools and platforms to reach them.

Case Study

OTT Walled Garden AdTech Stack — We designed a full-stack solution (self-serve platform, ad server, data lake, CDP, custom APIs, etc.). The client monetized first-party data and built a walled garden ecosystem.

This stack creates an infinity of possible use cases for our client.

Data Broker

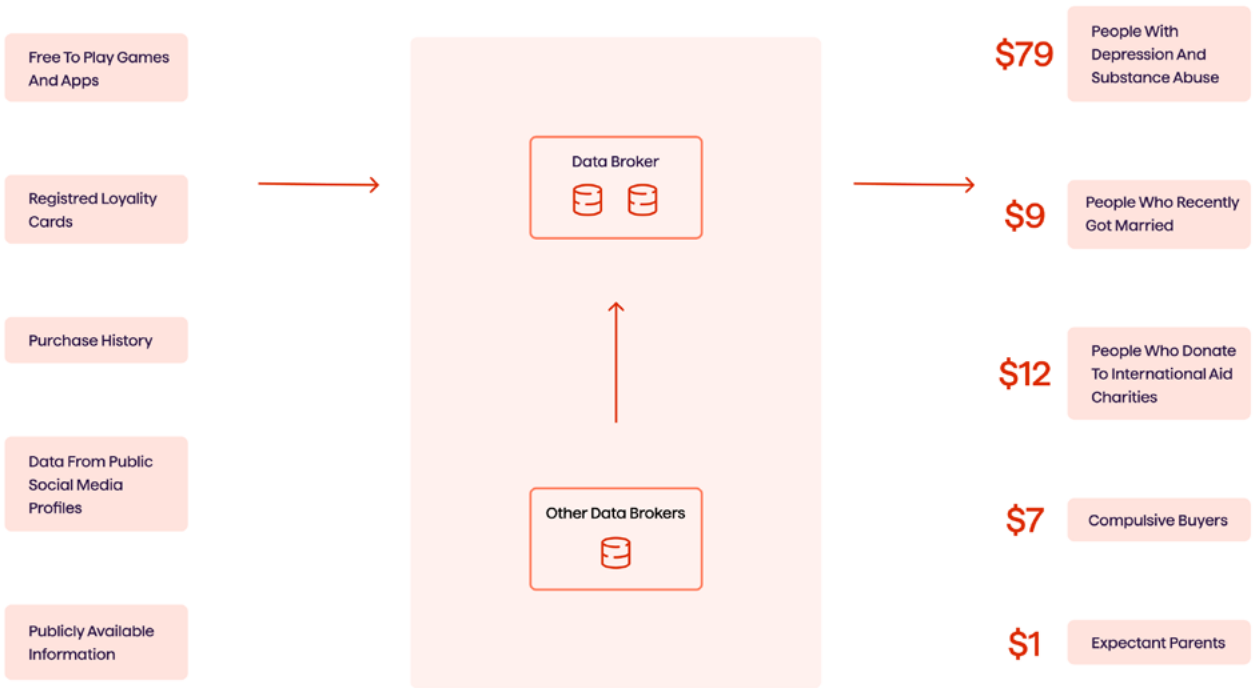

A data broker is a company that collects and aggregates personal information — such as income, ethnicity, political beliefs, or geolocation data — and then sells or licenses this data to third parties. Data brokers typically compile information from multiple sources, segment it into categories, and sell them to other companies for use in online advertising campaigns.
There are several types of data brokers that operate in certain industries:

- Marketing and advertising — Enhance ad targeting and campaign measurement.
- Identity verification and fraud detection — Help organizations such as banks confirm individual identities.
- People search — Collect publicly available information about individuals from social media and other online sources.

In the digital advertising and marketing industries, many DMPs act as data brokers and vice versa.

Advertising Agency

An advertising agency is a company that provides services to brands associated with creating, planning, and managing advertising campaigns.

Most agencies operate as independent, external firms serving a variety of clients, including businesses, multinational corporations, non-profit organizations, and governments.

Traditionally, brands hired ad agencies to produce television commercials and run print campaigns in magazines, newspapers, and on billboards, but also to take care of other forms of promotion and marketing.

With the rise of the internet, agencies now rely on a wide range of advertising and marketing technologies to design, run, and measure online campaigns. These are often referred to as interactive, creative, media, or digital agencies.

Agency Trading Desk (ATD)

An agency trading desk (ATD) is a unit within an advertising agency that provides programmatic managed services to brands.

ATDs handle programmatic media-buying activities on behalf of brands and consist of two layers: a services layer (media buyers, developers, account managers, etc.) and a technical layer (proprietary technology combined with external tools such as DSPs).

Most major advertising agencies operate their own trading desks, which manage programmatic media buying for the agency's clients.

AdOps

The AdOps department is responsible for setting up ad campaigns, trafficking tags, configuring header-bidding wrappers, and adjusting campaigns as needed. Both publishers and advertisers benefit from AdOps teams:

- Advertiser AdOps team — Configures campaigns in the advertiser's ad server and provides ad tags to the publisher.
- Publisher AdOps team — Sets up the campaign, adds the advertiser's ad tags to the publisher's ad server, and launches the campaign.

The terms *AdOps* and *ad trafficking* are often used interchangeably, but they have distinct meanings:

- Ad trafficking refers to the process of setting up, monitoring, and optimizing campaigns within ad servers and other AdTech platforms.
- AdOps refers to the people responsible for carrying out that process.

Verification and Measurement

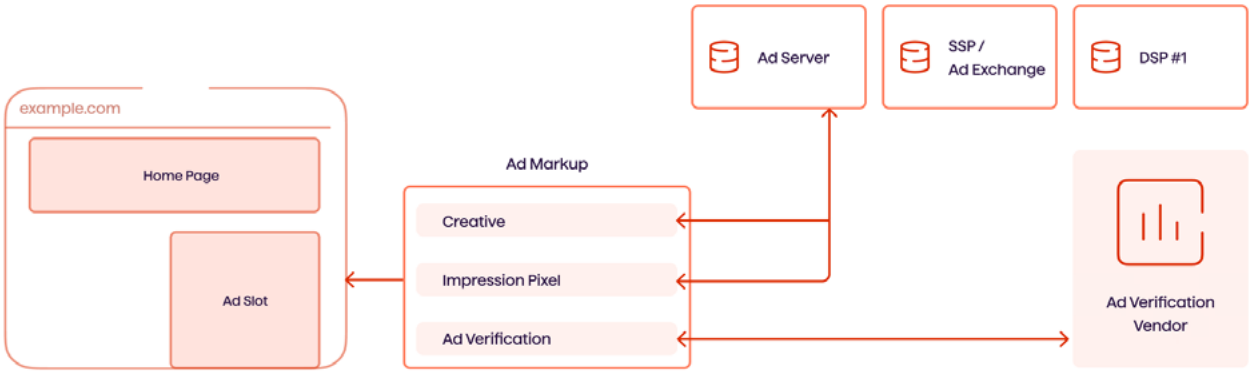

Verification services use technology to provide advertisers with additional information about their online advertising campaigns. The information provided by verification services can inform the advertiser about the following:

- Which websites the ads appeared on
- Where they were displayed (geolocation)
- What percentage of ads were viewable by the user
- Whether fraudulent traffic was detected (e.g., impressions, clicks, or conversions generated by bots)
- Whether ads appeared next to questionable content (e.g., illegal or offensive material)

Measurement and analytics companies provide detailed analysis of the performance and reach of online ads.

They also provide detailed insights into customer behavior, trends, and other user-centric data to help advertisers improve the performance and targeting capabilities of their campaigns.

It's important to note that the analytics mentioned here are different than the analytics used for tracking and reporting on website traffic or in-app user behavior.

Traditional web and app analytics belong to the MarTech ecosystem; however, many of these solutions are now expanding to include ad-campaign insights. This reflects a broader trend toward unified platforms that combine multiple functions and bridge both the marketing and advertising technology spaces.

Case Study

We developed two solutions: a simple analytics platform and an ad measurement and verification platform. For the latter, the central feature was the reporting interface, which included a homepage with several monitoring panels — covering reporting, viewability, brand safety and suitability, invalid traffic filtering, and geography.

This solution tracks more than 500M ads daily.

Dynamic Creative Optimization (DCO)

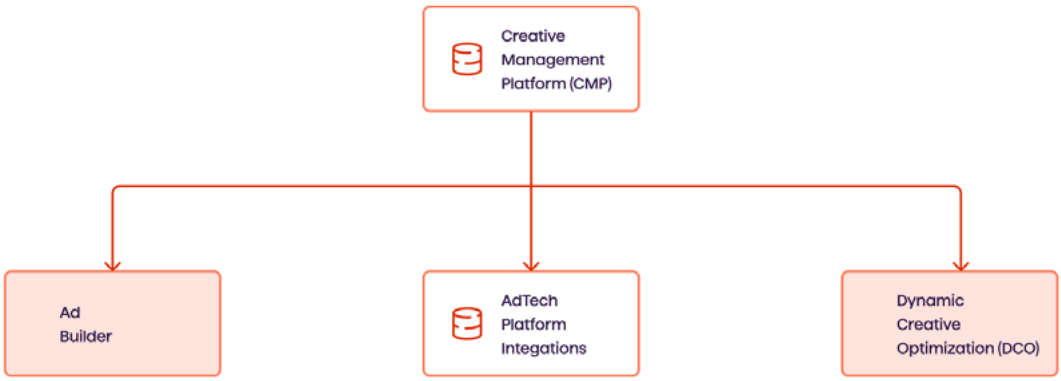

What Is a Dynamic Creative Optimization (DCO)?

Dynamic creative optimization (DCO) is a process whereby advertisers show personalized ads to individual users based on information known about them.

The software that powers this process is known as a dynamic creative optimization tool (or platform), which is essentially a piece of advertising technology that is responsible for creating, serving, and measuring these hyper-relevant advertisements to users.

With DCO, various ad components — such as backgrounds, text, images, value propositions, and calls-to-action (CTAs) — are dynamically adjusted to create the most personalized message for each user. These elements are tailored using the data associated with that specific user.

Within the DCO tool, advertisers utilize data about their clients, such as demographic data, geographic locations, interests, contextual

data, behavioral data, historical insights, and context-specific data. The DCO tool forms a new creative out of the provided information, tests it in real-time (A/B tests), and displays the best-performing creative at scale.

What sets DCO apart from other AdTech solutions is its real-time use of data to generate dynamic ads. By combining personalization and instant optimization, DCO significantly increases the likelihood that users will click and ultimately convert.

How Does a DCO Work?

To serve dynamically created and optimized ads, a DCO tool must integrate with existing AdTech platforms such as DSPs and ad exchanges. It then uses data feeds and machine learning algorithms to assemble personalized creatives.

To understand how a DCO tool works, it helps to see how it fits into the real-time bidding process:

1. User visit — When a user loads a website, the SSP sends a bid request to an ad exchange, which forwards it to all connected DSPs.
2. Bid evaluation — Each DSP evaluates the available user data against its campaign targeting criteria to decide whether to bid.
3. Bid response — The DSPs submit their bids, and the ad exchange selects the highest bidder. The winning ad is then cleared to display on the publisher's site.
4. Dynamic creative generation — Before the ad is shown, the DSP issues an ad call to the DCO. The DCO generates a hyper-relevant creative in real time and delivers it to the user.

Standardization

In the very early days of online advertising, it became apparent that there would be a need to provide a set of standards to ensure that different advertising technologies could communicate efficiently with one another and deliver ads in the correct format.

In 1996, the Interactive Advertising Bureau (IAB) was founded to standardize the online advertising industry. The IAB is responsible for developing technical standards, promoting best practices, conducting industry research, and educating companies on the value of digital advertising.

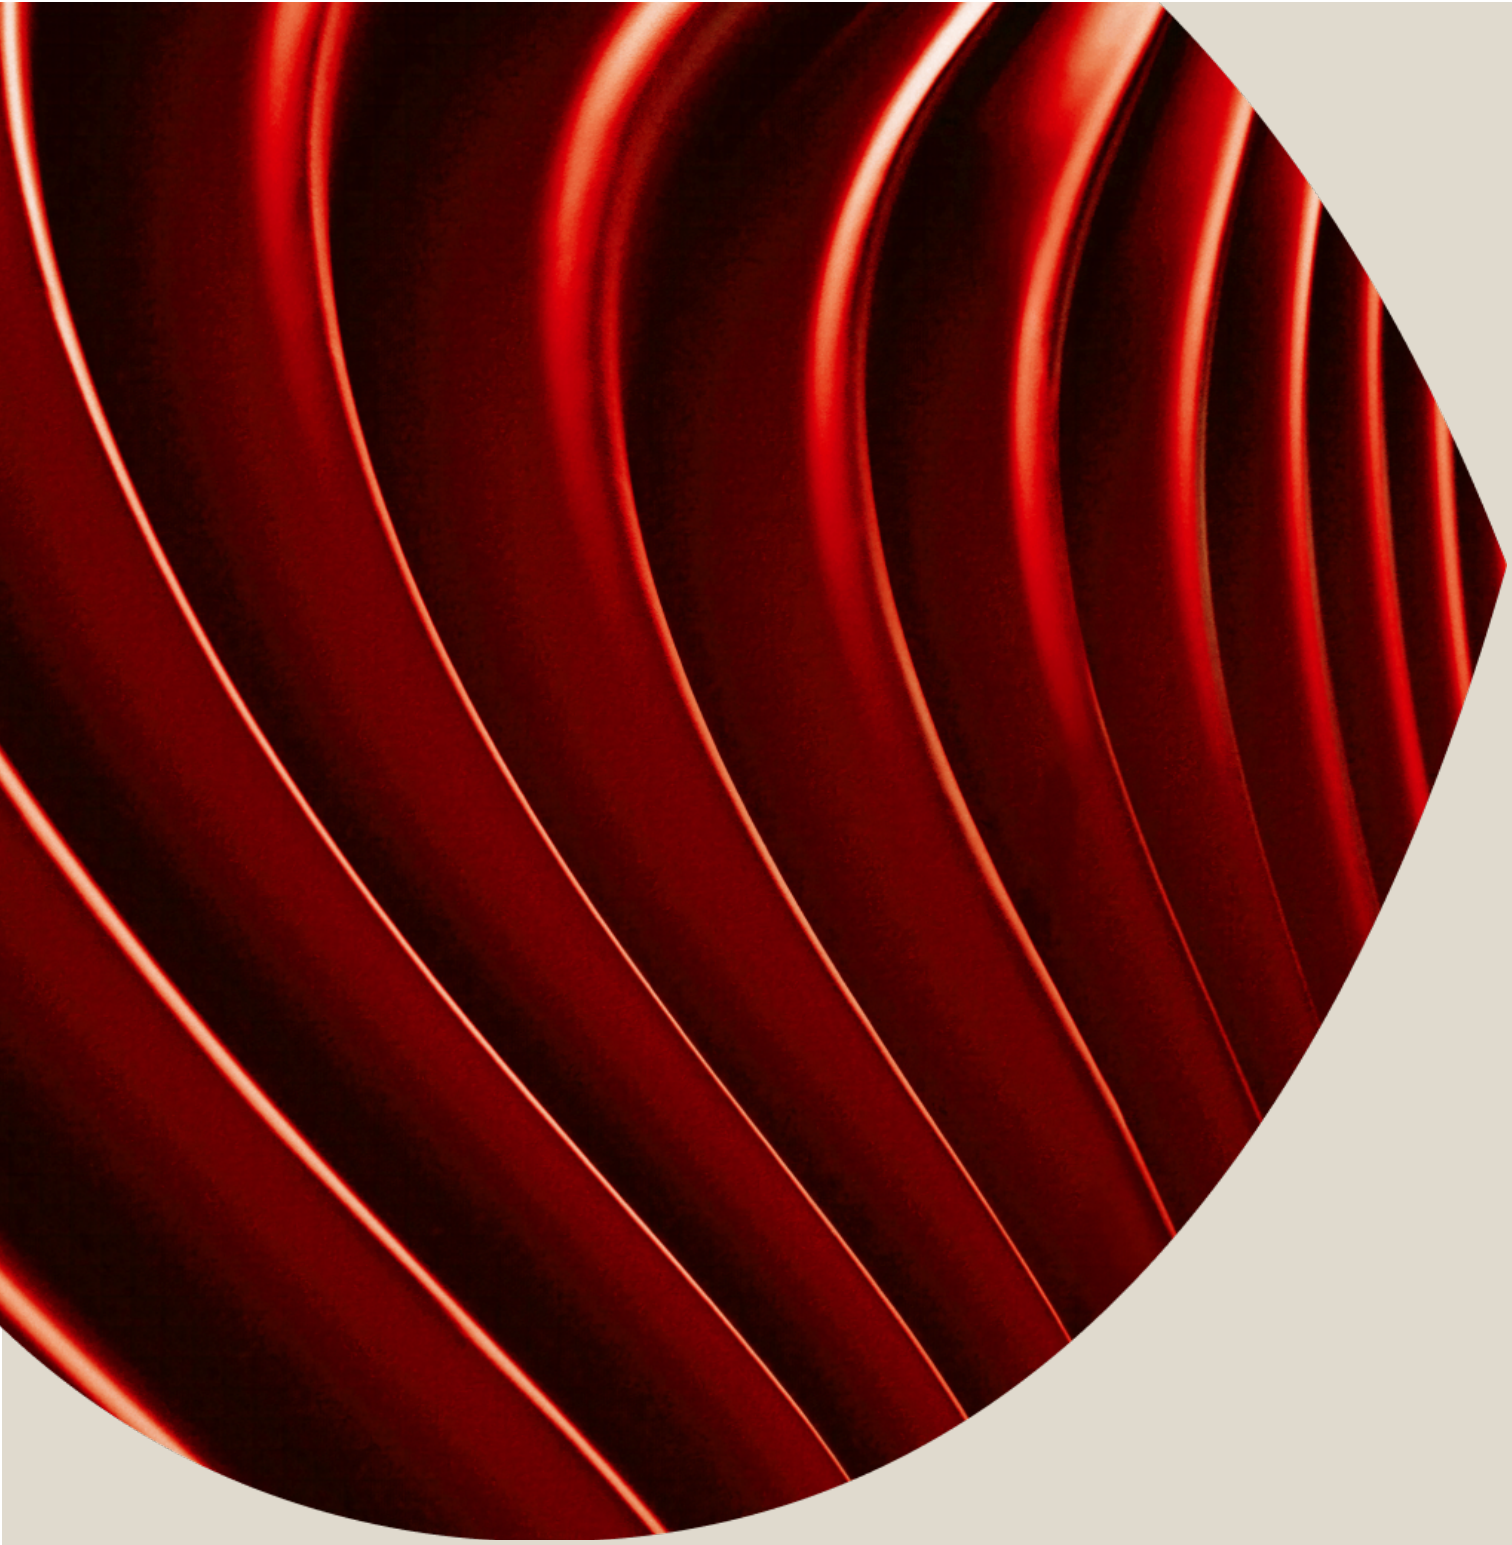

3. Advertising Mediums and Channels

With the rise of digital platforms, advertising has expanded far beyond traditional print and television.

This chapter examines the wide range of modern advertising formats — such as video, audio, and native ads — and distribution channels, including web, mobile, social media, and connected TV. It also highlights the growing role of AI in powering smarter, more targeted campaigns.

Mediums

A means of verbal or non-verbal communication. Examples include text ads, video ads, and radio ads.

Text and Image Ads

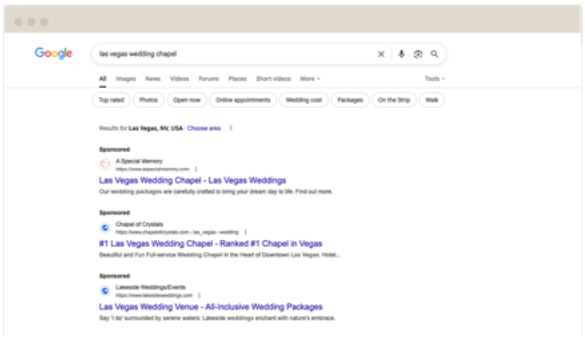

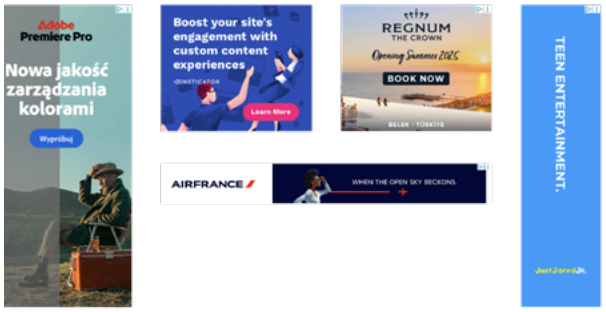

When online advertising first emerged in the late 1990s, the only available formats were text and image ads. Although many new formats have been introduced since then, text and image ads have remained consistently popular.

Native Ads

Native ads are designed to blend seamlessly with surrounding content by matching the design, format, and behavior of the webpage, application, or platform, creating a more natural user experience.

However, despite their integrated appearance, native ads must include a clear and prominent disclaimer indicating that they are paid advertisements.

Native ads are generally found on content-rich sites, such as news sites, blogs, and social networks. Here are some of the main formats:

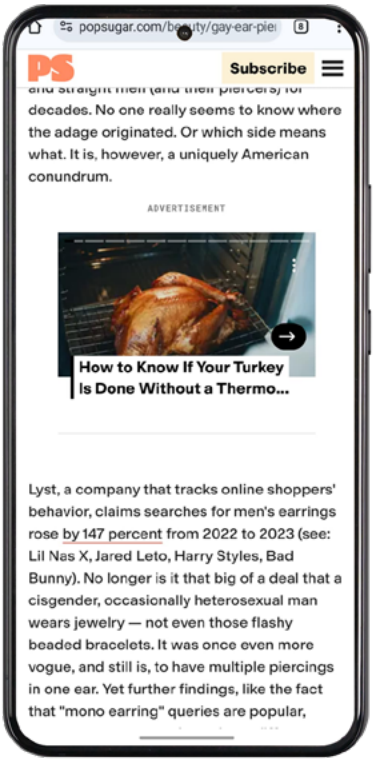

- **Content recommendation ads** — Perhaps the most subtle form of native advertising. These appear as suggested articles or links, often labeled with phrases such as *From the web* or *Recommended for you*. Because of their placement and style, they can easily be mistaken for actual site content.
- **In-feed or in-content ads** — Similar to content recommendation ads, these appear directly in article or content feeds (in-feed) or embedded within articles (in-content). They may take the form of text, images, or videos, blending naturally into the flow of the surrounding material.
- **Branded/ native content ads**— Unlike the first two formats, which usually direct users off-site when clicked, branded content is hosted directly on the publisher's site. Here, companies pay to feature content — such as sponsored articles or videos — presented in the same style as editorial pieces. These campaigns are often developed collaboratively between the brand and the publisher's editorial team, resulting in more engaging, high-quality content.

Video Ads

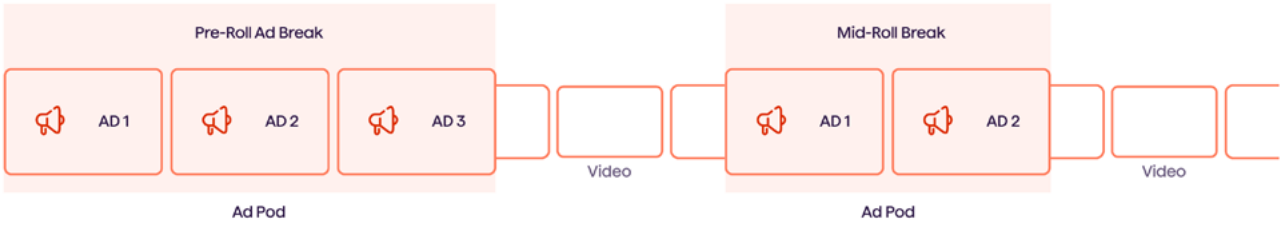

Video advertising refers to ads displayed in the form of a video. Depending on the channel, serving a video ad is similar to serving an image, text or native ad, except that the creative (i.e. the video) is sent to and displayed in a video player rather than as part of the web page itself.

Also, most video ads are served via protocols developed by the IAB Tech Lab:

- **VAST (Video Ad Serving Template)** — An XML framework that standardizes the delivery format for video ads across streaming platforms, ensuring compatibility between ad servers and video players.
- **SIMID (Safe Interactive Media Interface Definition)** — A standard designed to simplify ad interactivity by separating interactive features from measurement and verification, which are handled by other specifications.
- **VMAP (Video Multi-Ad Playlist)** — An XML framework that specifies where ads should be inserted within video content, enabling ad breaks such as pre-roll, mid-roll, and post-roll.

Rich Media Ads

Rich media is an interactive form of digital advertising that incorporates elements such as video, animations, audio, expandable content, and gamified features. These formats are designed to boost user engagement, increase click-through rates, and drive conversions.

Some examples of rich media ads are:

- Expandable ads
- Interstitials
- Video sliders
- AR/3D ads
- Gamified advertising
- Shoppable media

As ad platforms evolve, brands are introducing new formats such as QR codes and vertical-video experiences on TV screens, as well as AI-generated, immersive ads that seamlessly blend into the viewing experience — for example, pause-screen overlays with product placements.

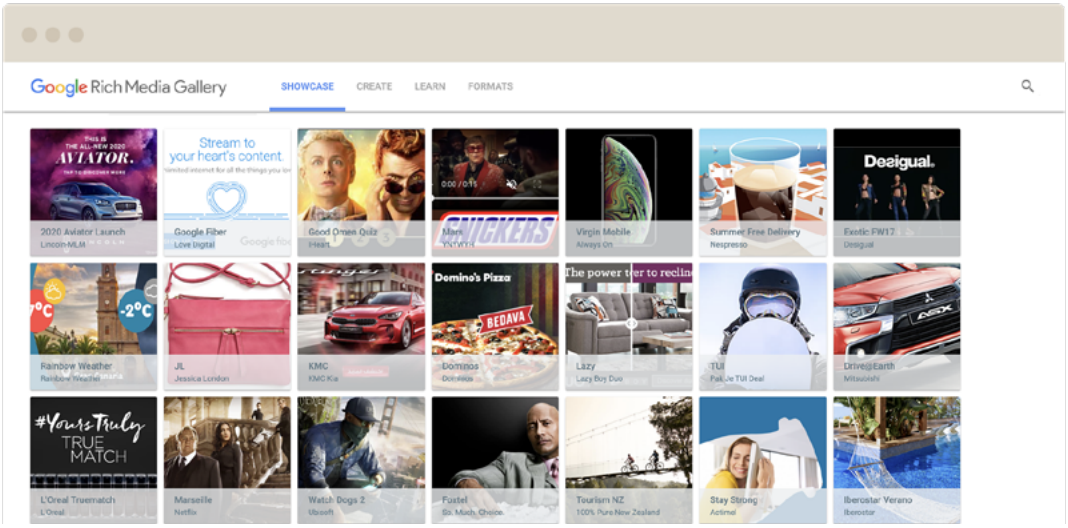

Audio Ads

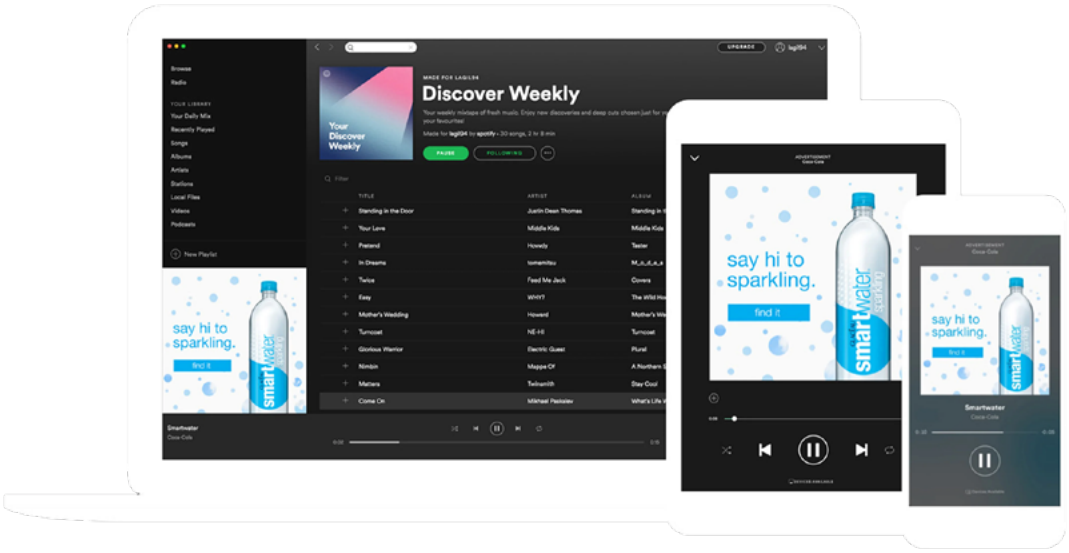

Audio ads are advertisements delivered in audio format, primarily on digital platforms like streaming services, podcasts, and online radio. They rely on sound to convey a message.
Due to the similarities between video and audio files, publishers and advertisers can use the latest version of VAST to serve ads and collect relevant data.

Audio ads also come in different formats:

- **Companion/banner ads** — Similar to display banners on websites and apps, these appear on the screen while a user listens to audio content (e.g., a podcast or music stream).
- **Ad pods** — Borrowed from video advertising; ad pods allow one or more audio ads to run within a single file. They may appear as pre-roll ads or mid-roll ads.
- **Dynamic ads** — Compared to static audio ads that are designed to announce the same message to the masses, dynamic ads change based on the information known about the user, their location, the time of day, and even the weather.

Channels

A means of transmission or distribution. Examples include display, social media, and TV advertising.

Web Advertising

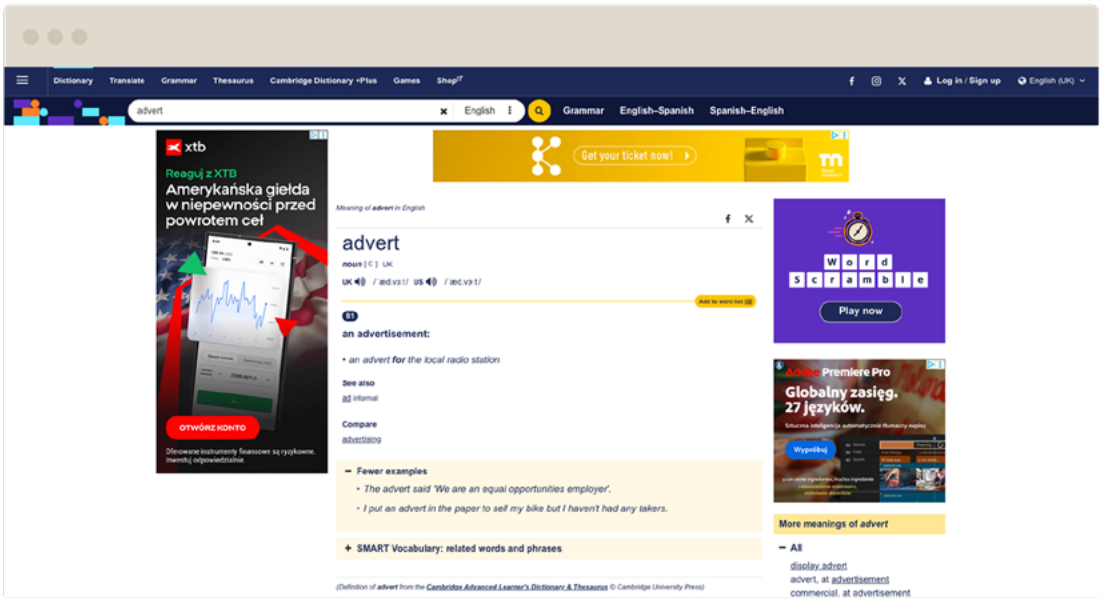

This channel refers to advertisements displayed in web browsers on desktops, laptops, and web browsers on mobile devices (smartphones and tablets).

Web ads — whether text, image, video, or rich media — are typically loaded through a small snippet of HTML (such as an <iframe> tag) or JavaScript (commonly called an ad tag). This snippet acts as both a placeholder and an instruction for the browser to fetch ad content from an ad server.

When a page loads, the browser often:

- Calls an ad server or ad exchange (via the snippet)
- Runs an auction in milliseconds to determine which advertiser wins the ad placement (RTB)
- Retrieves the winning ad creative (HTML, JavaScript, images, or video) from the advertiser's CDN or ad server. Rich media, CTV, and shoppable ads often require additional JavaScript SDKs for interactivity
- Renders the ad in an isolated container (often an <iframe>) to keep it separate from the page's own scripts
- Consent and privacy compliance (e.g., TCF v2.2 in the EU, CCPA in the US) can delay or block ad loading until the user agrees to tracking

Mobile App Advertising

Mobile app advertising, also known as in-app advertising, refers to the display of ads within mobile applications on smartphones, tablets, and other mobile devices.

Unlike mobile web advertising, which relies on a browser, in-app advertising requires the use of a software development kit (SDK) to render ads.

App developers integrate the AdTech vendor's SDK into their application, define the available ad space, and select both the ad medium (e.g., text, image, native, video) and ad format (e.g., interstitial, banner).

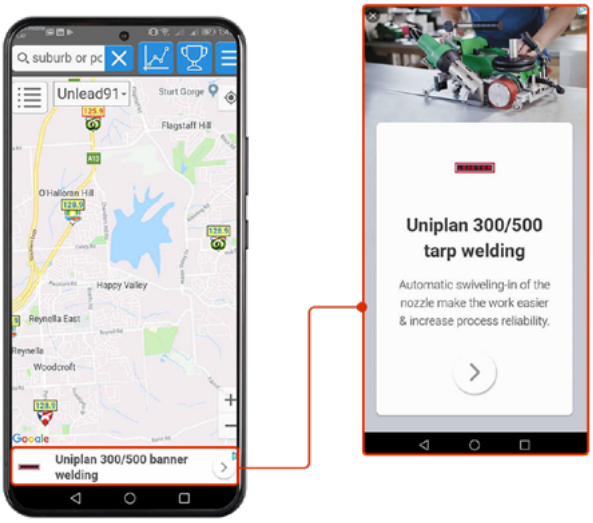

Social Media Advertising

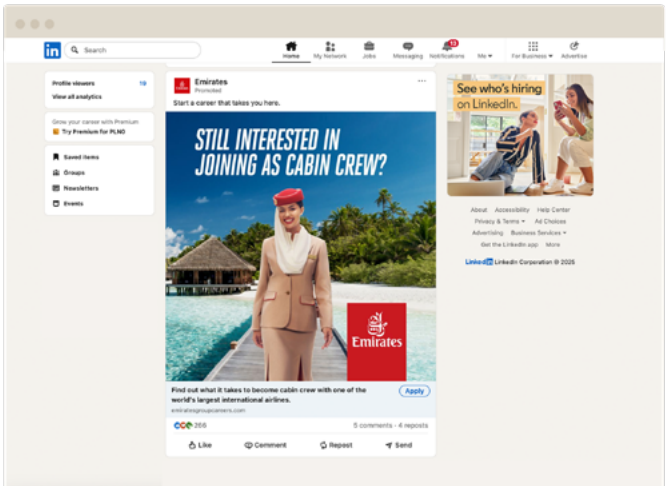

Although social media networks such as LinkedIn could be categorized under web or mobile-app advertising, they function as distinct advertising channels.

Many popular social media sites utilize native advertising, where the ads appear in and next to the news feed, making them look like organic content.

Advertising on social media offers several advantages for brands, advertisers, and agencies:

- **Retargeting** — Advertisers can retarget users using the email addresses they provided to a brand, often resulting in higher conversion rates.
- **Advanced targeting** — Platforms like Facebook and LinkedIn collect extensive demographic and behavioral data (e.g., name, age, location, interests, education), enabling highly refined audience targeting.
- **Cost-effectiveness** — Compared to other channels such as display or video, social media advertising can be more affordable while delivering strong reach and conversions.
- **Higher engagement** — Since most social ads are native in format, they typically generate more engagement (viewable impressions, clicks, likes, etc.) than standard banners and are less vulnerable to ad-blocking tools.

Advanced TV Advertising

TV advertising has long been the cornerstone of many brands and agencies. The ability to reach a large audience with an engaging and often entertaining message helps increase brand awareness and drive sales.

In the AdTech ecosystem, TV can be categorized into several forms:

- **Advanced TV** — Any form of TV other than the traditional broadcast, cable and satellite connections.
- **Over-the-top (OTT)** — Devices or services used to stream digital content to a connected TV. Examples of OTT services include Netflix, Hulu, Amazon Prime, etc.
- **Connected TV (CTV)** — Devices that connect to the Internet and allow viewers to watch video content from over-the-top streaming services. Examples of CTV include smartTVs, gaming consoles, and streaming devices. The IAB Tech Lab doesn't consider desktop computers, laptops, smartphones, and tablets as examples of CTV.
- **Addressable TV** — Aims to display different ads to different viewers during the same program. It can achieve this by using data collected via Internet Protocol TV (IPTV), which includes connected TVs and OTT devices and services, and set-top boxes.

These formats provide advertisers with advantages that traditional TV cannot match, such as precise targeting, improved attribution, and detailed measurement.

Digital-Out-Of-Home (DOOH) Advertising

Digital Out-of-Home (DOOH) advertising leverages digital screens — typically LED displays — positioned in high-traffic public spaces such as city streets, shopping malls, and transit hubs.

Unlike traditional static billboards, DOOH delivers dynamic content, including videos, animations, and interactive experiences. It also supports advanced targeting based on factors like location, time of day, and even weather conditions. With the added benefit of measurable performance data, advertisers can gain real-time insights into campaign effectiveness.

Common examples of DOOH include digital billboards in city centers, screens in bus shelters and train stations, interactive kiosks in shopping malls, and digital signage in airports.

Search Advertising

Search advertising is a form of digital marketing in which ads appear alongside search engine results when users enter specific keywords.

These ads are typically text-based and are triggered by user intent, making them highly relevant and timely.

Placed prominently on search engine results pages (SERPs), they help businesses reach potential customers at the exact moment they're actively seeking information, products, or services. These ads allow for precise targeting based on keywords, location, device, time of day, etc.

Search ads are offered both by search engines (e.g., Google) and by publishers, often in the form of promoted listings on e-commerce sites.

AI in Advertising Mediums and Channels

AI is transforming digital advertising mediums and channels by enabling smarter, faster, and more precise campaign execution across the entire advertising lifecycle.

- In display advertising, AI powers programmatic buying through real-time bidding (RTB), evaluating each impression based on user behavior, context, and predicted performance. DSPs use AI to automate bid strategies, optimize budget allocation, and select the most effective ad placements, ensuring maximum return on investment.
- For video and connected TV, AI enhances targeting and personalization by analyzing viewer habits, engagement patterns, and content preferences. It enables dynamic ad insertion and contextual targeting using machine learning models that process visual and audio cues to match ads with the most appropriate content environments.
- In search advertising, AI streamlines campaign management through automated bidding, keyword optimization, and predictive analytics. It helps advertisers reach users with high intent by forecasting click-through and conversion rates, adjusting bids in real time, and suggesting new keyword opportunities based on historical performance and trends.
- On social media platforms, AI enables hyper-personalized ad experiences by analyzing vast datasets on user interests, behaviors, and interactions. It drives dynamic creative optimization, where ad formats, visuals, and messaging are automatically adapted to each user's profile and real-time context, improving engagement and conversion rates.
- Mobile advertising benefits from AI through improved location-based targeting, app usage analysis, and in-app behavioral insights. AI helps deliver more relevant and timely ads while managing frequency capping and ensuring a seamless user experience.
- In programmatic advertising, AI connects all digital channels — display, video, mobile, audio, and native — into a unified buying strategy. It orchestrates cross-channel campaigns, ensures cohesive messaging, and provides real-time feedback for continuous optimization.

AI also plays a critical role in measurement and fraud prevention. It supports advanced attribution modeling, audience validation, and real-time anomaly detection to safeguard media investments and deliver more accurate insights into campaign performance.

4. Ad Serving

Ad servers are the backbone of digital advertising — determining which ad appears, where, when, and to whom.

This chapter breaks down how ad servers work, the differences between first-party and third-party servers, and the behind-the-scenes systems that power targeting, delivery, and tracking in milliseconds.

First-Party vs. Third-Party Ad Servers

The are two types of ad servers:

- **First-party ad servers** — Used by publishers to manage and optimize their own ad inventory.
- **Third-party ad servers** — Used by advertisers and agencies to serve, track, and measure ads across multiple publishers.

Although both types operate on similar technology, they fulfill different roles depending on whether they serve the needs of publishers or advertisers.

First-Party (Publisher's) Ad Server

A publisher's ad server is tasked with filling the ad slots on a website by matching ads from direct campaigns, real-time bidding (RTB) auctions, and other media-buying processes.

They do this by making decisions about which ads to show on a website based on the targeting parameters of campaigns set up by advertisers, serving those ads, and reporting on their performance.

Another feature of ad servers is inventory forecasting, which involves predicting how much inventory a publisher will have available in the future and the performance of ad campaigns based on current and historical data.

Third-Party Ad Servers

To obtain independent reporting for campaigns running across multiple publishers and ad networks, advertisers began using independent ad servers — now referred to as third-party ad servers (or advertiser ad servers).

A third-party ad server enables advertisers to:

- rack the performance (impressions, clicks, conversions, etc.) of the whole campaign across all publishers in a single system
- Measure overall campaign reach while accounting for co-viewership across publishers
- Verify the reports provided by the publishers

First-Party Ad Server (Publisher's Ad Server)	Third-Party Ad Server (Advertiser's Ad Server)
Enables AdOps teams to manage website ad slots, run multiple direct campaigns (i.e., direct deals between publishers and advertisers), and handle third-party tags from external ad servers or platforms such as SSPs.	Tracks campaign performance across all participating sites (e.g., reach, impressions, clicks, conversions), calculates return on investment (ROI), and attributes conversions to the appropriate publishers.
A first-party ad server helps publishers manage and forecast inventory fill rates across multiple advertisers, provides billing reports, and tracks both earnings and actual fill rates. It also evaluates the performance of third-party demand sources (e.g., via RTB) and direct deals, enabling publishers to see which sources purchase the most ad space and assign priorities accordingly.	Helps the advertiser optimize its future media buys based on past data, identify which sites and targeting criteria worked and which didn't, and run A/B tests to determine which ads deliver the best performance.
Enables publishers to forecast the volume of inventory that matches specific targeting criteria. For example, it can estimate how much traffic comes from New York State on a homepage ad placement, allowing the sales team to offer the appropriate quantity on insertion orders for advertisers.	Allows advertisers to audit and verify the numbers (i.e. performance metrics) for billing purposes.

How Does an Ad Server Work?

An ad server is a technology platform that automates the process of storing, selecting, delivering, and tracking digital ads. It ensures that the right ad is shown to the right user at the right time, while also collecting performance data for analysis.

The ad serving process typically unfolds in four stages:

Step 1: Ad Request Initiation

When a user visits a website or opens an app containing ad space, the site sends an ad request to the ad server. The request contains contextual and technical information, such as user's location, device type and browser, content of the page, cookie or user ID data for targeting.

Step 2: Ad Selection Logic

The ad server applies business rules and targeting criteria to review available ad campaigns, apply segmentation filters (e.g., demographics, geography, interests), evaluate campaign goals and priorities (e.g., CPM, CPC) and select the best-fitting ad for delivery, based on performance and eligibility.

Step 3: Ad Delivery

Once the ad is selected, the ad creative (HTML, image, video, etc.) is served in the designated ad slot. Media assets may be pulled from a Content Delivery Network (CDN), and the ad is displayed seamlessly to the user within milliseconds.

Step 4: Tracking and Analytics

The ad server continuously logs impressions, clicks, conversions and viewability, and engagement metrics. This data is compiled into reports for advertisers and publishers, providing insights to optimize performance and maximize ROI.

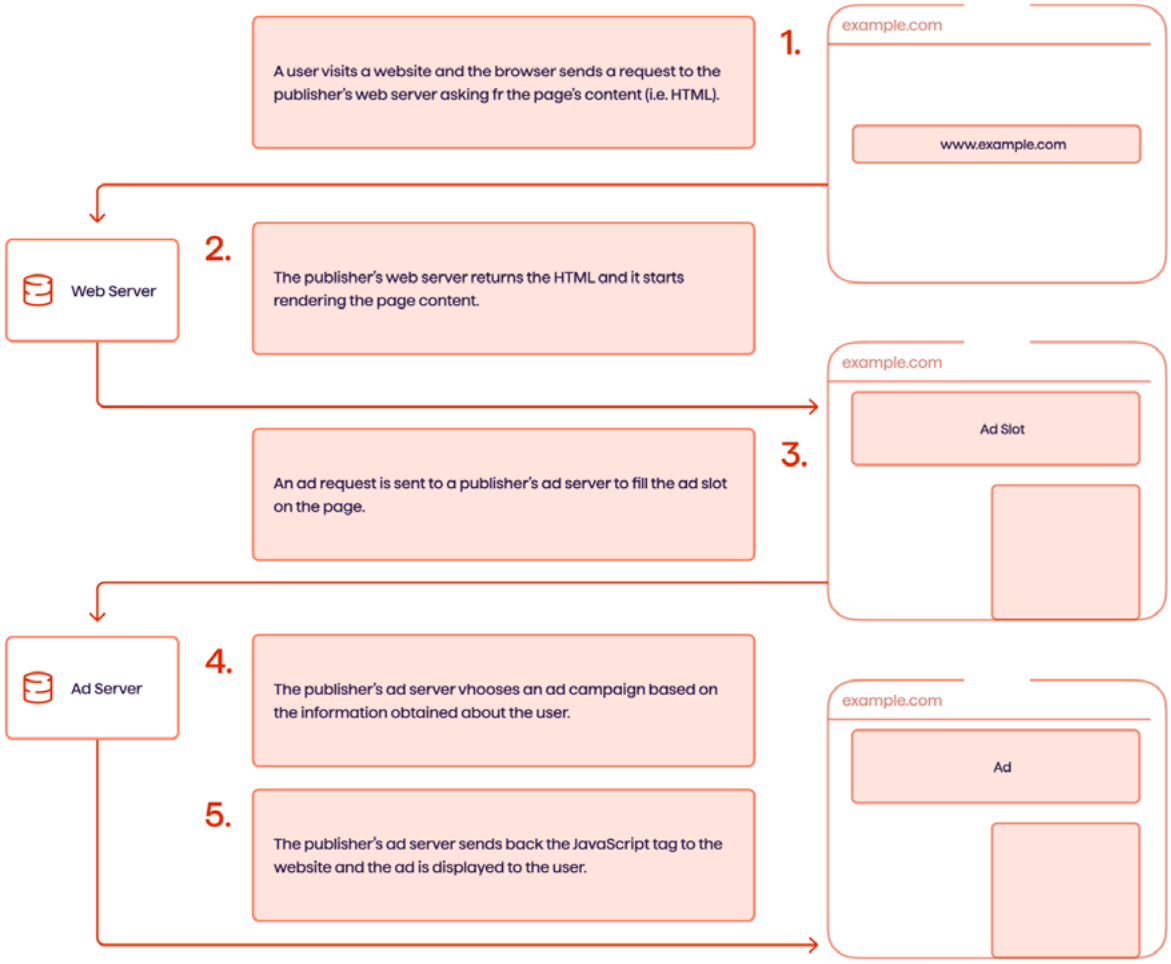

In programmatic advertising environments, the ad serving process includes a real-tim auction:

- The ad request is sent to an SSP
- The SSP communicates with DSPs
- Advertisers bid in real time for the impression via Real-Time Bidding (RTB)
- The highest bidder wins, and their ad is served through the ad server

Below are two detailed technical examples of how ad serving works. The first illustrates how ad serving works with just a publisher's ad server, and the second illustrates the ad-serving process with both a publisher's and an advertiser's ad server:

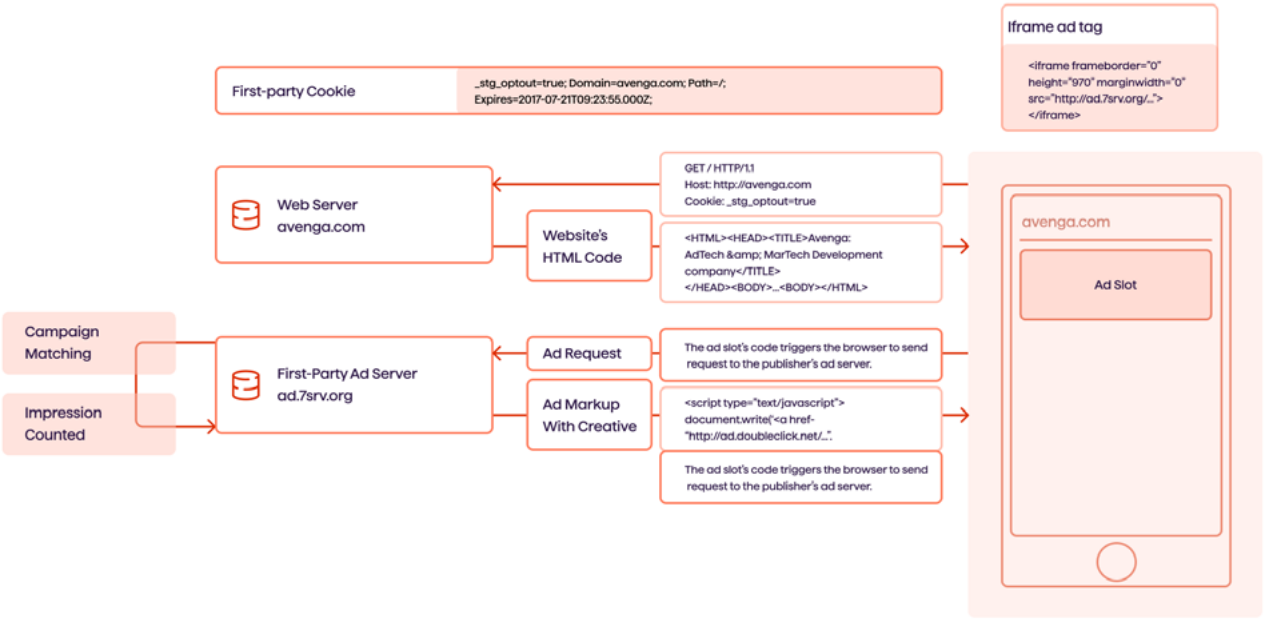

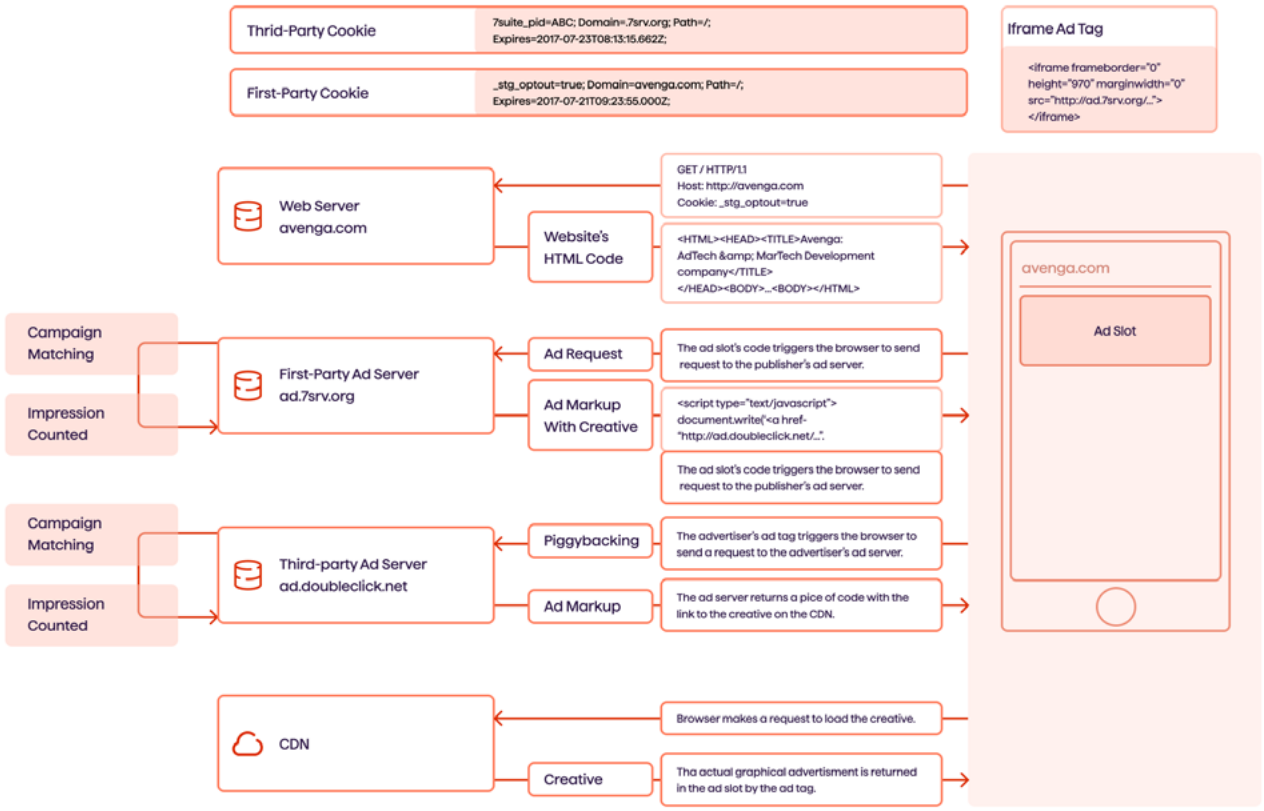

The Anatomy of an Ad Server

An ad server runs on a set of tightly integrated systems. Each has a specific job, but together, they power the entire ad delivery process.

The following sections outline the key ad-serving components that work behind the scenes to deliver the digital ads users encounter every day.

In simple terms, an ad server functions like a high-speed dispatch center — constantly scanning the digital landscape, making split-second decisions, and delivering tailored ads to the right audience at the right time.

1. The Ad Decision Engine

This is the decision-maker — the component that determines which ad gets shown. It processes millions of signals in real time:

- Who is the user?
- What page are they on?
- What targeting rules apply?
- Which campaigns have budget left?

It balances logic and strategy to deliver the most relevant, best-performing ad without violating campaign rules or budgets.

2. Inventory Management System

Imagine you're hosting a conference with hundreds of speakers and limited stage time. The inventory system manages the "stages" — the ad spaces across websites, apps, and devices. It:

- Tracks what inventory is available
- Organizes placements by format, location, and priority
- Matches them to campaigns that qualify to fill those slots

Without this system, no one would know where or when ads should appear.

Case Study

Personalized Messaging Tool — We developed a platform that allows merchants to create personalized campaigns, enabling targeted onsite messaging to promote their core business offers.

Our solution can render personalized messages within 50 milliseconds.

3. Creative Repository

Every ad needs a home before it's served. That's what the creative repository is for. It securely stores:

- HTML5 banners
- Video ads
- Image assets
- Third-party tags

The ad server pulls from this vault to instantly serve creatives — optimized for format, size, and speed.

4. Targeting and Segmentation Module

This is where the magic of personalization happens. The targeting engine reads data from cookies, device IDs, geolocation, or user segments and matches it to campaign criteria. Whether it's "women aged 25–34 in New York" or "returning users who abandoned cart last week," this module makes it possible.

5. Delivery Engine

The delivery engine takes the selected ad and gets it to the user in milliseconds. It connects with CDNs (Content Delivery Networks) to reduce load times and handles everything from rendering banners to auto playing videos.

6. Tracking and Measurement System

This part uses invisible tracking pixels, SDKs, or beacons to collect this data, which then feeds into dashboards for real-time insights.
The tracking system logs:

- Impressions
- Clicks
- Viewability
- Engagements
- Conversions

7. Reporting and Analytics Interface

This module transforms raw data into actionable insights, so marketers can make better decisions, faster. Advertisers and publishers can log in to see:

- Campaign performance
- Budget pacing
- Engagement rates
- Recommendations for optimization

8. Identity and Frequency Control Layer

Nobody likes seeing the same ad 100 times.
The control layer enforces:

- Frequency caps (e.g., 3 impressions per user per day)
- User ID matching for cross-device delivery
- Sequenced messaging (showing Story A before Story B)

9. API and Integration Layer

Modern ad servers don't live in isolation.
They need to play nice with dozens of systems such as DSPs, CRMs, analytics platforms, consent managers, and more. This layer makes it all possible, providing real-time connectivity, customization, and automation.
A programmatic environment or large-scale publishers' ad servers may also include:

- **Real-time bidding engine** — Auctions impressions in milliseconds to the highest bidder
- **Header bidding manager** — Lets multiple demand partners bid before the ad server chooses
- **Brand safety and verification tools** — Protect brands from showing ads on inappropriate or fraudulent content

There are even more ad-serving components:

- **Ad tags** — Is a piece of code that is inserted into an ad slot to display an ad. There are different types of ad tags depending on their implementation, for instance, JavaScript ad tags, iframe ad tags, SafeFrame, img ad tags, video ad tags, etc.
- **Ad markup** — Is a piece of code that's retrieved from an ad server or some other AdTech platform via an ad tag and rendered in an ad slot. Ad markup is responsible for two main activities:

1. Loading the actual creative file into the ad slot
2. Tracking the impression by loading tracking tags (pixels) for measurement, ad verification, viewability, etc.

5. Ad Targeting and Budget Control

Ad targeting ensures ads reach the right audience, while budget control helps advertisers spend wisely.

This chapter examines the main methods of audience targeting — ranging from contextual and behavioral approaches to device-based and demographic filters — and explains how advertisers manage campaign budgets through capping, pacing, and frequency limits to maximize efficiency and minimize waste.

Contextual Targeting

Contextual targeting displays ads that are relevant to a website's content rather than relying on visitor data. This approach has long been common in traditional media, such as magazines and newspapers, and continues to be effective online.

Advertisers and publishers often use contextual targeting on its own or alongside other targeting methods, as it is particularly effective for content tied to specific topics or themes.

How Does Contextual Targeting Work?

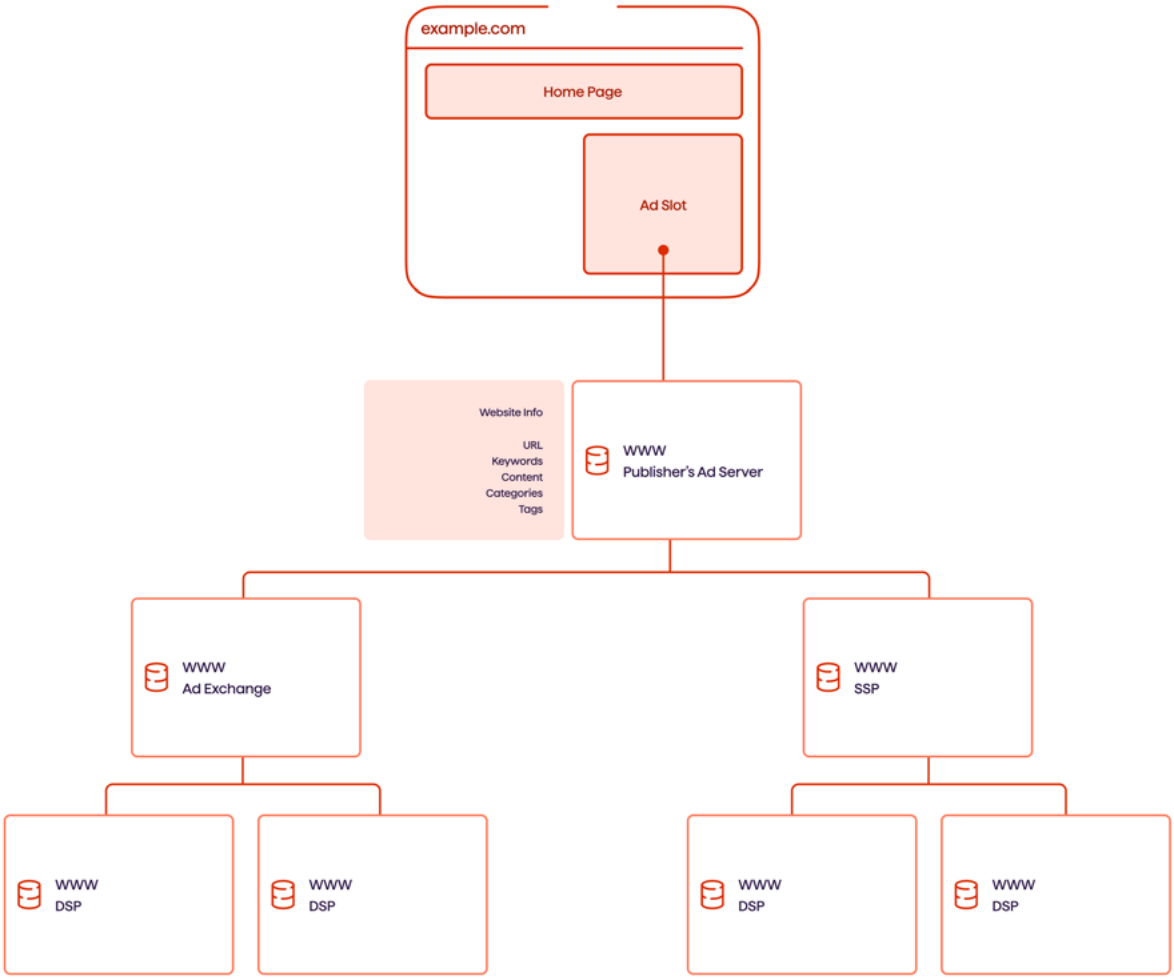

A step-by-step explanation of how contextual targeting works:

1. A web crawler scans URLs and categorizes the content and ad placements.
2. When a visitor accesses a website, information associated with the URL is passed to the ad server via the ad request.
3. This ad request and the contextual information is passed on to other AdTech platforms, such as ad exchanges and supply-side platforms.
4. The ad exchanges and SSPs relay this information to demand-side platforms, which then bid on the impression.
5. The winning DSP then sends the ad back to the publisher and displays the ad to the visitor.

Benefits of Contextual Targeting

Although contextual targeting may appear simpler than other methods, it offers several important benefits for advertisers and publishers:

- Most contextual ads don't rely on personal data, helping advertisers and publishers reduce their exposure to privacy and data-protection regulations, like the GDPR
- Contextual ads can offer safer brand protection
- They are proven to increase purchase intent
- They are found to be less unnerving than behaviorally targeted ads, while still being based on the interests of users — e.g. advertisers can show ads for smartphone plans to people reading articles about smartphones

Keywords

Keywords are valuable targeting variables for advertisers as they provide an accurate picture of the type of content the website visitor is consuming.
There are a couple of ways an ad server can identify which keywords are on a given web page, such as via the tags used on the page by the editor that highlight the key topics covered in the content and via extracting the keywords found on the page, usually done either with JavaScript or server-side web crawling.
These keywords are then typically passed to the ad tag, so that the ad server receives them in the ad request and uses them in the decisioning process (when deciding which ad to serve).
So, for example, if an advertiser wants to display ads promoting a new smartphone, they will target visitors who read articles containing the keywords smartphone, mobile phone, etc.

Ad Slot and Ad Position

Advertisers may also choose to display ads of a specific size or in a particular location on a webpage — for instance, a 728×90 px banner positioned at the top of the page.

Since this form of targeting is fairly broad, it's typically used in combination with other targeting methods.

Publisher's URL

Targeting users based on the publisher's URL is very similar to the way advertisers target consumers in print media.

By displaying ads on specific websites (specific domains, URLs, section of a website), advertisers can reach a broad range of online consumers based on their interests, rather than based on their demographic information like age and gender.

IP and Geolocation

Geolocation targeting involves displaying ads to users based on their current location. For example, if an online consumer was reading news articles on a laptop in Chicago, they could very well see ads promoting shops, restaurants and services in the Chicago area.

But how do advertisers know where consumers are located? When a user's browser sends an ad request, it includes the device's IP address. The ad server uses an external database to map that IP address to a specific geographic location — typically the country, region, or city.

Native mobile apps can provide even more precise data by passing a smartphone's exact longitude and latitude coordinates from its GPS to the ad server.

In such cases, the ad server is capable of targeting users within a radius of a certain point (e.g., within five miles of a given point of sale). This makes geolocation targeting much more accurate and precise than the IP address method.

However, even GPS data can occasionally be imprecise or manipulated. To improve reliability, data providers often aggregate and cross-reference multiple data points and sensors.

Browser Type, Operating System and Device Type

Every ad request to the ad server carries a user-agent HTTP header, e.g.:

```
User-agent: Mozilla/5.0 (iPhone;
CPU iPhone OS 18_0 like Mac OS X)
AppleWebKit/605.1.15 (KHTML, like Gecko)
Version/18.0 Mobile/22A5297f Safari/605.1
```

In the example above, 22A5297f identifies as an iPhone 15 Pro Max, and we see it's running iOS 18.0.

From this data, advertisers can determine the user's operating system, browser type and version, and device type, brand, and model (for example, identifying an iPhone or Android smartphone).

Targeting online consumers based on their hardware or software enables advertisers to reach a specific audience with a highly relevant message.

For example, a mobile-gaming company could display ads promoting its newly developed Android game to consumers using Android-powered smartphones and tablets.

IAB Content Taxonomy

The IAB provides a standard for the categorization of websites, called content taxonomy. Advertisers are able to purchase digital ad space based on the categories supplied in the ad request and can also choose not to show their ads on websites based on the categories.

Day of Week and Time of Day

Displaying ads to consumers based on the day of the week and even time of day can allow advertisers to reach their desired audience at the right time and also avoid wasting their ad budget.

For example, an advertiser that works for a large pizza restaurant could advertise its Friday night specials on Friday afternoons between 3-8PM.

Similarly, if a brand observes that engagement peaks during certain hours, it can schedule ads to run during those windows. This approach not only improves targeting efficiency but also boosts engagement, click-through rates, and conversions.

Behavioral Targeting

Behavioral targeting, also known as online behavioral advertising (OBA), enables advertisers and publishers to display relevant ads and marketing messages to users based on their web-browsing behavior.

The types of data collected for behavioral targeting include:

- Pages viewed
- Previous search terms
- Amount of time spent on a website
- Ads and buttons clicked
- Content viewed and downloaded
- Purchases
- Date of the last website visit
- Other information about their interactions with various websites

1. Data Collection

Advertisers, publishers, and data management platforms collect data about the actions users carry out across different websites. This data is often referred to as event data and includes page views, products views, products purchased, and other interactions on a website or mobile app. This data is then tied together via identifiers stored inside third-party and first-party cookies in web browsers, or mobile IDs in mobile apps.

User profiles are then created to consolidate each user's event data and link future activity to that same profile. Identifiers such as cookie IDs or mobile device IDs connect a user's behavior across multiple sites, ensuring that new event data is assigned to the correct profile.

2. Audience Creation

Advertisers and publishers then create audiences that are made up of individual user profiles.

For example, an advertiser could create an audience that includes people who have viewed a given product more than three times a month, signed up for a newsletter, and have visited their website at least 15 times in the past 60 days.

3. Application of Data

The advertiser then uses those audiences for ad targeting in its online media campaigns, resulting in the ads being more relevant to the users and increasing the chances of them converting (e.g. purchasing a product).

Benefits of Behavioral Targeting

The vast amount of data available to marketers enables them to build highly detailed user profiles and deliver ads tailored to each audience segment.

The core idea behind behavioral targeting is mutual benefit — users see ads that match their interests, while websites improve engagement and overall user experience.

However, online users have become aware of how online advertising companies collect and use their data, which has resulted in some users becoming concerned about this type of targeting and resulted in the rise of ad-blocking software.

Retargeting

Retargeting is the practice of displaying ads to users who have previously interacted with a brand.

For example, if an online visitor views a pair of shoes, then they'll likely see the same pair of shoes advertised on different websites.

This process works by embedding a 1×1 transparent image, known as a pixel, on a webpage. When the page loads, the pixel sends a request to an AdTech platform (such as a demand-side platform). As the image is returned to the browser, the DSP places a cookie on the visitor's device (if one doesn't already exist).

Later, when that visitor browses another website, the DSP recognizes them via the cookie and displays a retargeted ad — for instance, the shoes they previously viewed.

Here's a simplified example of how retargeting works:

1. An online shopper visits blueshoes.com and views a pair of shoes.
2. The retargeting service code located between the <footer> tags send a request for a 1×1 pixel.
3. The retargeting service sends back the 1×1 pixel and saves a cookie to the shopper under their domain (ads.retargetser.com), storing information about the shopper and their behavior, such as the product they viewed.
4. The shopper leaves blueshoes.com and visits a different website — news.com — and sees an ad for the exact same pair of shoes that they were looking at previously.

The DSP is able to identify the same visitors across different websites by syncing cookies with other AdTech platforms — for example, supply-side platforms and ad exchanges.

Demographic

Demographic targeting is one of the most powerful and precise forms of audience targeting. It's often used alongside other methods to further refine campaign reach. However, it can be challenging because most publishers don't collect demographic data directly from their visitors — with major exceptions like Facebook and Google, which gather this information through user accounts.

Examples of demographic information include, but are not limited to:

- Age
- Gender
- Annual income
- Marital status
- Parental status
- Occupation

Case Study

Our AdTech and MarTech experts rebuilt an email retargeting tool designed to help eCommerce stores recover lost revenue from high cart-abandonment rates.

The solution we created for our client has generated over $56 million in revenue.

For example, an advertiser promoting baby products could target female users aged 20-40 who have one or more children. By combining demographic filters with other targeting methods, the advertiser can ensure ads reach the most relevant audience.

When running campaigns on independent AdTech platforms (i.e., not through Google, Facebook, or LinkedIn), advertisers typically apply demographic targeting using audience segments from a DMP or via demographic data contained in the User object of OpenRTB bid requests.

When running campaigns using demographic data on Facebook and other platforms that collect this type of data directly from users, advertisers simply need to set up the targeting criteria inside those platforms. data on Google and Facebook and other platforms that collect this type of data directly from users, advertisers simply need to set up the targeting criteria inside those platforms.

Controlling a Campaign's Budget

When running an ad campaign, advertisers focus not only on improving performance — by refining messaging, placement, and targeting — but also on reducing ad waste.

Ad waste refers to paid impressions that fail to reach the intended audience due to factors such as fraud, low viewability, or mistargeting. By minimizing waste, advertisers can optimize their budgets and improve overall efficiency.

Below are several key areas advertisers monitor to control and manage campaign spending — all of which can be configured within an ad server or demand-side platform.

Budget Capping

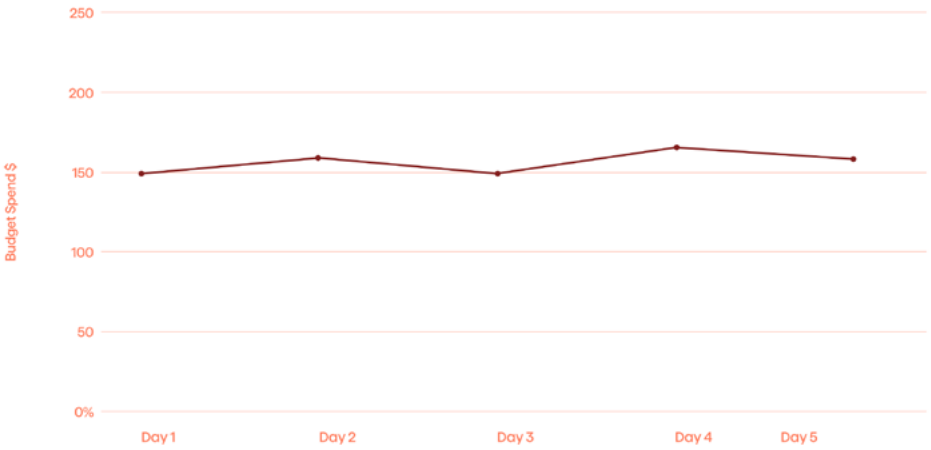

Adding a budget cap to an ad campaign sets a limit on how much money will be spent.

For example, you could set a daily budget cap of $150, meaning that once the campaign has spent that amount of money, no more ads will be shown that day.

Limiting a campaign's budget involves putting a total and daily cap on the amount of money a campaign spends. In this case, once the daily budget has been spent, the campaign will stop for that day.
Some platforms will add a certain percentage (e.g., 20%) to the daily budget to help advertisers get the most out of their campaign.

While increasing the daily cap can help compensate for under-delivery on certain days, it can also lead to premature budget depletion.

Budget Distribution (aka Pacing)

Pacing refers to how quickly or slowly the budget of the campaign is spent, which impacts how many impressions are served to online consumers over any given period during the campaign's lifetime.

For example, you may want to get your ads in front of as many consumers as quickly as you can, so you would choose to deliver the maximum number of impressions as soon as possible, known as ASAP pacing.

Alternatively, you may choose to distribute your ads evenly and allow them to run for the duration of the campaign's proposed dates, known as uniform pacing.
Examples of ASAP and uniform pacing:

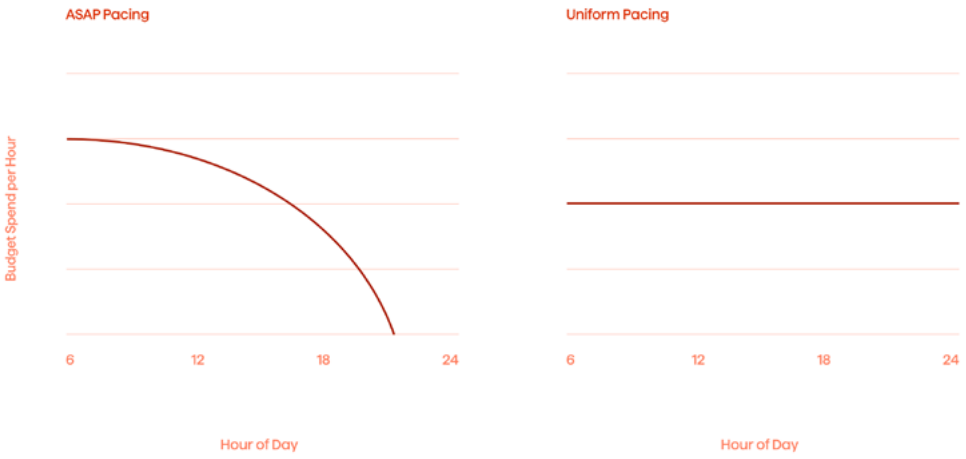

In practice, pacing must also account for fluctuations in ad traffic throughout the day and the availability of impressions that meet the campaign's targeting criteria.

Ad platforms dynamically adjust how quickly they spend the campaign budget based on factors such as performance, traffic volume, and the cost of inventory or audience segments:

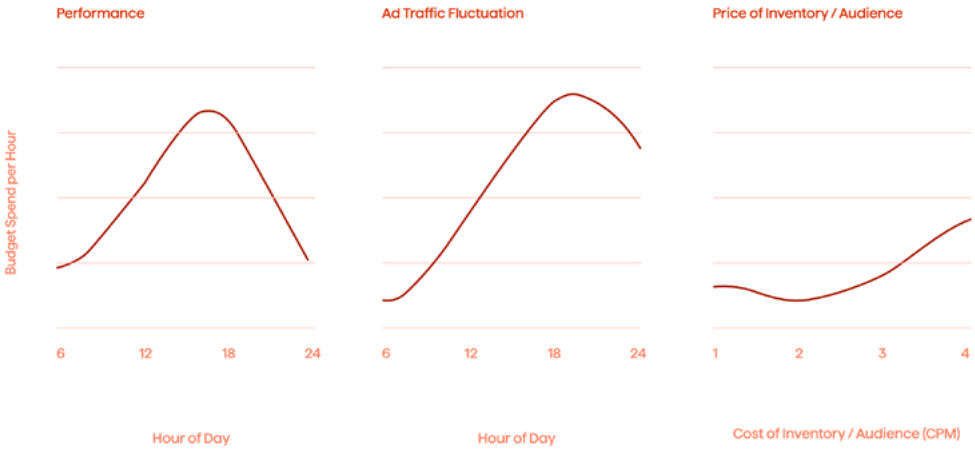

Case Study

Ad Banker – we created a DSP component that enables several bidders to receive and respond to bid requests simultaneously while preventing budget overspending. The ad banker is a part of our AdTech Foundations - a collection of essential components for building robust AdTech and MarTech platforms.

Ad Banker reduces budget overspend for 36% - from 37.2% to just 1.2% million in revenue

Frequency Capping

Frequency capping involves limiting the number of times the same ad is shown to a given visitor. For example: three impressions per visitor per 24 hours.
Frequency capping is important because it:

- Limits budget waste
- Helps to improve a campaign's overall reach
- Prevents so-called "overexposure" (a user's frustration over seeing the ad multiple times)

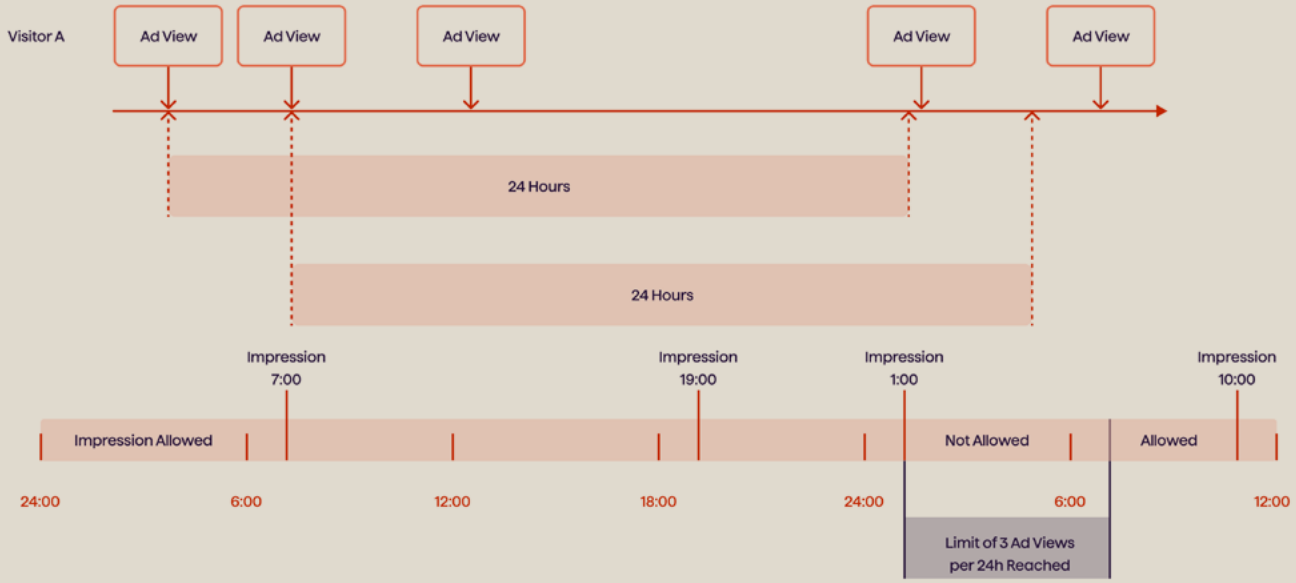

Each time frequency capping is evaluated, the system counts the number of impressions served within the defined time window. For example, with a 24-hour cap of three impressions, the ads could be shown at 7:00, 19:00, and 1:00 the following day — after which the user wouldn't see the same ad again until the next 24-hour cycle begins.

6. Tracking and Reporting

How could one assess digital advertising campaign performance without tracking and reporting solutions?

This chapter explains how impressions, clicks, and conversions are recorded, the technologies behind them, and how reports turn raw data into actionable insights — ensuring campaigns are optimized, accountable, and transparent.

How Impression, Click and Conversion Tracking Works

Impression Tracking

An impression is recorded each time an ad is displayed to a user. For example, if a user visits a webpage, views an ad, and then refreshes the page to see the same ad again, that counts as two impressions.

The most common way to track impressions is through a 1×1 transparent image — often referred to as an impression tracker or impression pixel. This image is returned by the ad server within the ad markup, and it's only loaded when the browser renders the ad — ensuring that impressions are counted when the ad is actually shown, not just when selected.

Example of an impression tracker (Google Ad Manager):

```
<IMG SRC="https://ad.doubleclick.
net/ddm/trackimp/Nxxxx.site-keyname/
Byyyyyyy.n;dc_trk_aid={ad_id};dc_trk_
cid={creative_id};ord=[timestamp];dc_
lat=N;dc_rdid=Czzzz;tag_for_child_
directed_treatment=I?" BORDER="0"
HEIGHT="1" WIDTH="1" ALT="Advertisement">
```

Click Tracking

Clicks are typically tracked via click trackers — special redirect URLs provided by ad servers. When a user clicks an ad, the tracker counts the click and then redirects them to the landing page.

To allow both the publisher and advertiser to track clicks, a redirect chain is created. The user is first sent through the publisher's click tracker, then through the advertiser's click tracker, and finally to the landing page. Additional redirects may be inserted for purposes like ad verification or intermediary ad network tracking.

Example (Google Ad Manager):

```
http://pubads.g.doubleclick.net/
gampad/clk?id=123456789&iu=/1234/
adunit&t=page%3Dsports
```

Click URL Macros

Click trackers often support macros to dynamically pass the next redirect URL. For example:

```
https://ad.example.org/click?ad_
id=123456&redir_url=%%click_url%%
```

This allows the ad server to expand the macro (%CLICK_URL%) into the next URL in the redirect chain.

Conversion Tracking

A conversion occurs when a user completes a predefined goal (e.g., purchase, form submission, download). Conversion tracking is essential for evaluating campaign performance and is especially important in cost-per-action (CPA) models.

There are two types of conversions:

- **Click-through conversion (CTC)** — The user clicks an ad and later converts
- **View-through conversion (VTC)** — The user sees (but doesn't click) an ad and converts later

CTC rate = (Conversions ÷ Clicks) × 100
VTC rate = (Conversions ÷ Impressions) × 100

Attribution Windows

This refers to the timeframe in which a conversion can be credited to an ad interaction (e.g., 24 hours to 30 days). Choosing the right window is crucial, as too long may misattribute conversions and too short may miss valid ones.

Conversion Tracking Methods

- **Pixel method** — A conversion pixel fires on the success page, linking the event to the cookie ID from the original click or view
- **Server-side method** — The cookie ID or click ID is passed via server-to-server calls — especially useful for affiliate marketing and mobile app installs where pixels can't fire

Additional Pixels and Piggybacking

When tracking impressions, additional third-party tracking pixels may be embedded in the ad markup to report impressions across multiple platforms. This is referred to as piggybacking.

Piggybacking enables external platforms, such as measurement tools and ad verification services, to log the same impression event by triggering their trackers simultaneously with the primary ad server's pixel.

JavaScript and Server-Side Tracking

JavaScript-Based Tracking

This method is commonly used on web pages to load and trigger tracking pixels when users interact with ads. It relies on browser-side tracking, making it effective for capturing real-time engagement data — such as mouse movements, scroll depth, and ad visibility — which are often used for viewability measurement.

Server-Side Tracking

This method eliminates dependency on browser events. Instead of firing pixels, user identifiers (e.g., cookie IDs or click IDs) are transmitted directly to the ad server through backend communication. Common use cases include:

- Affiliate marketing — Ensures accurate CPA attribution
- Mobile app installs — Uses SDKs and postback URLs to confirm successful installs

Viewability and Verification

Viewability refers to whether an ad was actually seen by a user. For example, an ad may technically be “served” but never viewed if it appeared below the fold and wasn’t scrolled into view.

Verification platforms ensure that:

- Ads are viewable
- Fraudulent activity (e.g., bots) is detected
- The environment is brand-safe

These platforms use JavaScript tags, tracking pixels, and sometimes server-side validation to confirm impression quality and delivery accuracy.

Two Common Methods for Counting Impressions

Dual Pixel Method

The publisher’s ad server embeds two 1×1 tracking pixels in the ad markup — one for the publisher and one for the advertiser. Both pixels fire simultaneously, allowing each party to log the same impression independently.

Request-Based Method

In this approach, each ad server — the publisher’s and the advertiser’s — records an impression when it receives an ad request. However, this method can create discrepancies if the ad markup fails to render, such as in cases of network latency, browser errors, or ad blockers.

Reporting

The reporting function of an AdTech platform provides critical visibility into campaign performance by displaying a variety of metrics and dimensions. Reports benefit both advertisers and publishers by surfacing insights about user interactions, campaign delivery, and revenue generation.

For every impression, click, and conversion event, the AdTech platform stores data attributes such as:

- Timestamp (impression, click, conversion)
- IP address
- Campaign ID, Line Item ID, Creative ID
- Geolocation, Browser, OS
- Publisher domain and placement

...and others.

Metrics

A metric is a calculated value that helps quantify campaign performance. Many are derived from raw data or other metrics, such as:

- **Impressions** — Number of times an ad is displayed
- **Clicks** — Number of times an ad is clicked
- **Conversions** — Number of desired user actions (e.g., purchases or sign-ups)
- **Reach** — Number of unique visitors or devices reached
- **CTR (Click-Through Rate)** — (clicks ÷ impressions) × 100
- **CPM (Cost Per Mille)** — Cost per 1,000 impressions
- **CPC (Cost Per Click)** — Cost per click
- **CPA (Cost Per Action)** — Cost per conversion
- **CVR (Conversion Rate)** — (conversions ÷ clicks) × 100
- **Amount spent** — Total media cost
- **Revenue** — Total conversion value

Example comparison:

Campaign	Pricing Model	Impressions	Clicks	Conversions	Cost	eCPM	eCPC	eCPA
#1	CPM	1,000,000	1,500	10	$4,000	$4	$2.67	$400
#2	CPC	1,000,000	2,000	50	$10,000	$10	$5	$200
#3	CPA	1,000,000	2,500	80	$15,000	$15	$6	$187.50

Effective Cost Per Mille, Click, and Action (eCPM, eCPC, eCPA)

"Effective" metrics (e-prefixed) standardize campaign performance across pricing models. These help compare campaigns that may be priced differently (CPM, CPC, CPA).
Formulas:

- eCPC = total spend ÷ number of clicks
- eCPA = total spend ÷ number of conversions
- eCPM = (total spend ÷ impressions) × 1,000

These metrics also allow publishers to retroactively evaluate what their revenue would've looked like under different pricing models.

Return on Investment (ROI)

To evaluate efficiency:
ROI = (revenue − spend) ÷ spend × 100%

Campaign	Revenue	Spend	ROI
1	$1,200	$1,000	+20%
2	$900	$1,000	-10%

Note: ROI doesn't include operational costs — important for physical goods, but less relevant for digital products with low marginal costs.

Dimension and Subdimensions

Dimensions are attributes used to break down and analyze data. Common examples include:

- Country
- Device type
- Browser
- Time of day
- Campaign
- Line Item
- Creative
- Publisher domain
- OS
- OS version
- Geolocation

Subdimensions (or drill-downs) allow deeper breakdowns
(e.g., country → carrier → line item → ad).

Case Study

AdOps Reporting Dashboard — We built a dashboard to collect, aggregate, normalize, process, and visualize data from an ad mediator platform. The goal here was to view campaign performance metrics and manage and control publishers' (our client's clients) inventory. With our solution, publishers gained the possibility to filter reports by demand partners, devices, and time intervals.

Three groups are the end users to the solution: AdOps, system users, employees.

Filtering

Filtering, also known as segmentation, narrows down the dataset used in reports to focus on specific criteria or dimensions.

Common filters include:

- Date range
- Campaign hierarchy
  (advertiser → IO → line item → ad)
- Geographic or technical filters
  (e.g., country = Poland OR Germany)

You can also apply "include/exclude" filters for more precise segmentation.

Technical Considerations of Reporting

There are several technical variables that can affect the accuracy and interpretation of reporting data. Below we describe the most common considerations.

Delays

Reports often lag behind real-time events. While approximated data might be available within minutes, accurate data suitable for billing typically takes longer (up to 24 hours).

Reporting Time Zone

If two systems operate in different time zones, their reports will not align. It's essential to confirm and standardize reporting time zones when comparing data across platforms to ensure accuracy.

Data Retention

To manage data volume efficiently, many platforms reduce either retention duration or data granularity over time:

- Last month — Hourly granularity
- 1-12 months — Daily granularity
- 1 year — Campaign-level summaries only

Discrepancies: Trust, But Verify

Discrepancies — differences in reported metrics between systems — are a well-known and often sensitive issue in AdTech. They typically occur between publisher and advertiser reports and can directly affect billing accuracy and overall trust.

Most discrepancies stem from technical or implementation issues, especially when using client-side tracking (browser or SDK).

Discrepancy Tolerance

Discrepancy = (publisher impressions − advertiser impressions) ÷ advertiser impressions × 100

IAB guidelines recommend tolerating discrepancies up to 10% (based on the advertiser's data). If the discrepancy is within this range, the publisher's data is typically accepted for billing. For instance, 7% discrepancy is acceptable, but 15% discrepancy requires investigation.

Reconciliation

Reconciliation is the process of comparing datasets from multiple systems to resolve reporting inconsistencies. In AdOps, this is often a manual process that involves:

- Pulling reports from both the publisher and advertiser systems
- Matching campaign IDs, timestamps, and event types
- Identifying where impression, click, or conversion counts diverge or conversions diverged

Common Causes of Discrepancies

1. Human and Implementation Errors
    - Incorrect or partial pixel placement
    - Misconfigured macros or missing cache busters
    - Differences in campaign start/end dates across systems
2. Configuration Differences
    - Mismatched reporting time zones
    - Different fraud filters, traffic-validation criteria, or viewability rules
    - Varying methodologies for counting impressions (e.g., pixel-fired vs. server-served)
3. Client-Side Tracking Limitations
    - Poor connectivity or latency preventing pixel load
    - JavaScript errors or browsers restrictions blocking scripts
    - Browser settings blocking scripts or pixels
    - URL length limitations truncating redirect paths
    - Creative file size affecting load speed
    - Resource-heavy pages delaying pixel fires

7. Media-Buying

Media buying transforms advertising budgets into a measurable audience reach. Whether executed through automated programmatic auctions or direct deals, it represents the point where strategy becomes execution.

This chapter explores the ecosystems, technologies, and deal types that connect advertisers with publishers, along with the metrics and optimization methods that ensure every impression counts.

Programmatic Media Buying

In the past, running a digital ad campaign was a manual process. Advertisers and publishers negotiated deals directly, exchanged ad tags via email, and coordinated placements and schedules by hand. This worked on a small scale but quickly became inefficient and error prone as digital advertising expanded.

Programmatic media buying introduced speed, scalability, and precision to the ad-buying process. At its core, programmatic means using software, data, and algorithms to automate the buying and selling of digital ads.

Instead of coordinating every detail manually, programmatic platforms allow advertisers to set up rules (like audience targeting or bid limits), and the system takes care of the rest — executing the campaign, adjusting bids, and optimizing performance, often in real time.

While humans are still involved in setting strategy and monitoring performance, much of the day-to-day execution is handled by tech. This shift gave rise to optimization practices — both manual and automated — that continuously refine campaign outcomes.

Manual Campaign Optimization

Manual optimization involves a member of the advertiser's team reviewing campaign data and making strategic adjustments to improve performance.

For instance, if a campaign is active across five countries and one underperforms, a media buyer might lower bids in that region or exclude it entirely. They could also test new creatives or adjust the ad schedule to boost engagement.

Typical tasks in manual optimization include:

- Adjusting CPC, CPM, or CPA bids based on performance
- Filtering out low-performing segments (e.g., underperforming publishers, devices, or geos)
- Testing different ad creatives (A/B testing)
- Reviewing campaign breakdowns to identify patterns and trends
- Personalizing creatives for specific audiences (e.g., geo-targeted messages)
- Experimenting with new platforms or inventory sources

This process gives advertisers a lot of control but requires time and constant monitoring. It's especially effective in smaller or highly strategic campaigns where nuance and manual judgment matter.

Automated Campaign Optimization

As digital campaigns scaled and data volumes grew, manually optimizing every campaign became unsustainable. That's where automated optimization shines.

In automated optimization, the system uses historical and real-time data to make performance-based decisions automatically. This can include things like adjusting bids per impression, reallocating budget to better-performing segments, or serving the highest-performing creative — without any human involvement.

These algorithms are often built into DSPs or other AdTech platforms and can optimize based on:

- Conversion probability
- Click-through rate
- Engagement level
- Cost-efficiency (e.g., reducing eCPA)

However, automation is only as good as the data that fuels it. Without sufficient, high-quality data, algorithms can misinterpret trends or fail to optimize effectively. For this reason, automated optimization performs best after a campaign has accumulated enough performance history to train and refine its decision-making models.

Manual vs. Automated Optimization: What's the Difference?

Here's a quick comparison of the two approaches:

	Manual Optimization	Automated Optimization
Control	Full control over every change	Limited, algorithm decides based on set parameters
Speed	Slower, changes applied manually	Fast, adjustments happen in real time
Scalability	Harder to scale, requires human input for each change	Easy to scale, works well across many campaigns
Flexibility	High, can make strategic and nuanced decisions	Lower, follows algorithmic logic
Responsiveness	Delayed response to performance trends	Instant reaction to performance changes
Data Dependence	Can work with smaller data sets and human interpretation	Requires a large amount of high-quality data to be effective
Use Case	Best for small campaigns, tests, or sensitive brand campaigns	Best for large-scale or always-on campaigns
Human Effort	High, requires daily monitoring and tweaking	Low, monitored but not constantly adjusted by humans

Most modern campaigns use a hybrid approach leveraging automation for scale and speed, while relying on manual input for strategy, troubleshooting, and fine-tuning.

For example, a DSP might automatically adjust bids and pacing in real time, while the advertiser manually refines audience targeting or creative messaging to align with broader brand objectives or context-specific insights that algorithms might overlook.

Programmatic Direct

Programmatic direct (known also as programmatic guaranteed, programmatic reserved, or automated guaranteed) is a method of buying and selling digital advertising inventory that combines the control of traditional direct deals with the automation of programmatic technology.

Unlike real-time bidding (RTB), programmatic direct doesn't involve auctions. Instead, media buyers and publishers negotiate terms directly — such as placement, volume, flight dates, and price (typically a fixed CPM), and then use AdTech platforms to execute and manage the campaign automatically.

How It Works?

1. Advertisers browse digital media catalogs
2. They select placements, dates, and impression volumes
3. Creatives and tracking pixels are configured
4. The order is placed via a platform
5. The publisher audits the campaign for compliance
6. Once approved, the campaign runs without manual trafficking by AdOps teams

Advantages

- Guaranteed inventory — Buyers lock in premium ad placements in advance
- Predictability — Fixed pricing and delivery commitments provide transparency and stability
- Efficiency — Automation reduces manual coordination for both advertisers and publishers

Disadvantages

- Limited targeting — Typically restricted to contextual and placement-level targeting
- Higher CPMs — Premium guaranteed placements come at a higher cost
- Potential underutilization — Publishers risk unfilled inventory if direct deals don't sell out

Real-Time Bidding (RTB)

Real-time bidding (RTB) transformed digital advertising by enabling advertisers to purchase individual ad impressions in real time instead of buying bulk inventory in advance.

Emerging in the late 2000s, RTB was developed to address inefficiencies in early ad networks that often over-delivered or under-delivered campaigns. As publishers began connecting with multiple ad networks, latency and coordination issues became increasingly problematic.

To solve this, supply-side platforms (SSPs) — initially known as network optimizers — and demand-side platforms (DSPs) were introduced to automate and balance the supply-demand relationship. Soon after, ad exchanges took the concept further by introducing impression-level, real-time auctions.

RTB operates much like a financial stock exchange — but instead of trading shares, participants bid on digital ad impressions. Each time a user loads a webpage, a lightning-fast auction (usually within 100 milliseconds) determines which advertiser's creative will appear.

What Is RTB and How Does It Work?

Real-time bidding (RTB) is a standardized protocol that allows advertisers to bid on each impression in real time, based on user and contextual data. Initially adopted to monetize remnant inventory, RTB has since evolved to include premium placements across display, video, mobile, and connected TV.

Key characteristics of RTB include:

- Data-driven bidding — Decisions are based on user demographics, device type, geolocation, and behavioral data
- Efficiency — Advertisers pay only for impressions that match their targeting criteria
- Scale — A single platform can provide access to thousands of publishers and billions of daily impressions

What Is the RTB Project (OpenRTB)?

The RTB Project, formerly known as the *OpenRTB Consortium* and now referred to as OpenRTB, is a group led by the Interactive Advertising Bureau, consisting of AdTech companies from both the demand and supply sides.

Launched in November 2010, the RTB Project defines a common API specification that enables seamless, automated communication between platforms for buying and selling digital media. In other words, it provides the shared "language" that DSPs, SSPs, and ad exchanges use to trade ad impressions in real time.

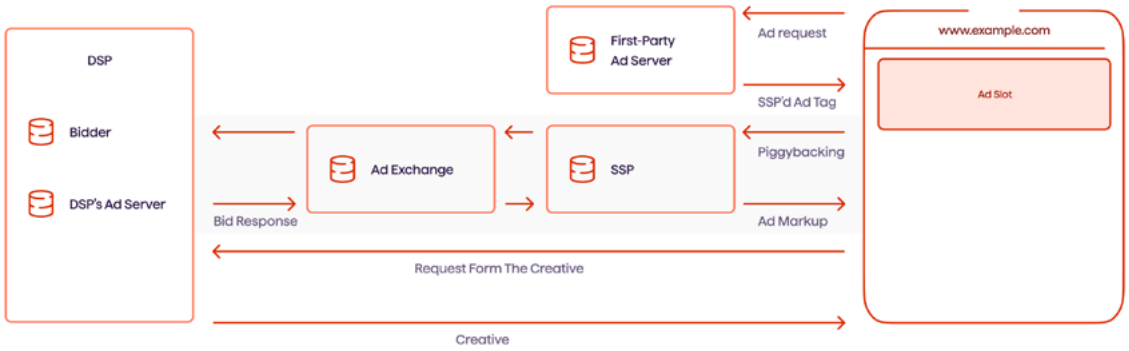

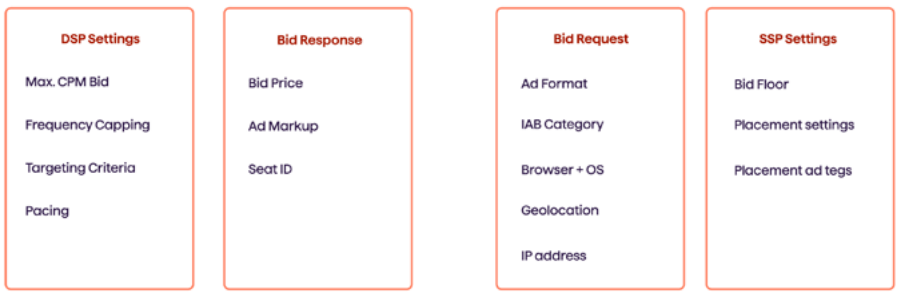

Here's how RTB works:

1. User visits a webpage
2. The page sends an ad request to the publisher's first-party ad server
3. If no direct campaigns match, an SSP tag is triggered
4. The SSP issues a bid request to ad exchanges with user and page data
5. Bidders (DSPs) evaluate the request and return a bid response
6. The ad exchange selects the highest bid (often second-price model) and notifies the winning DSP
7. The DSP's ad markup is sent to the user's browser
8. The browser retrieves the ad creative from the DSP's ad server or CDN
9. The ad displays, and impression tracking pixels are fired

This entire process takes 100–150 milliseconds — faster than the blink of an eye.

What Do Bid Requests and Bid Responses Contain?

Real-time bidding (RTB) transactions are conducted through JSON-based bid requests and responses that adhere to the OpenRTB specification.
Each message contains a structured set of objects describing the ad opportunity and the environment in which it will appear.

Common OpenRTB object types include:

- **Impression (Imp)** — Contains details such as ID, ad type (banner, video, native), and dimensions
- **Device** — Describes the user's device, including user agent, geolocation, operating system, and make/model
- **User** — Includes identifiers, gender, and year of birth
- **Geo** — Provides location data such as latitude, longitude, country, and region
- **Publisher** — Identifies the publisher's ID, domain, and name, as well as optional data segments like behavioral or demographic classifications

These objects provide DSPs with the necessary context to determine bid value. These objects provide DSPs with the necessary context to determine bid value.

Benefits of RTB for Advertisers

RTB allows advertisers to track campaign performance in real time and make immediate adjustments to improve results.
It also enables more precise audience targeting by leveraging both first-party and third-party data. Additionally, many RTB platforms include built-in fraud detection tools, helping advertisers reduce wasted spending on invalid traffic.

Benefits of RTB for Publishers

Publishers benefit from RTB through increased revenue, as competitive bidding often drives up CPMs. They also gain access to a wider pool of potential buyers, which enhances fill rates. Furthermore, RTB supports dynamic floor pricing, allowing publishers to adjust minimum prices in response to real-time demand, maximizing the value of their inventory.

Private Marketplace (PMP)

PMPs are invite-only RTB environments where premium publishers offer exclusive access to top-tier inventory. Advertisers receive Deal IDs that allow them to bid before inventory hits the open market.

Advantages for Advertisers

- Early access to premium, brand-safe inventory
- Improved viewability and transparency into placement quality
- Greater control over where their ads appear

Advantages for Publishers

- Higher CPMs compared to open exchanges
- Selective access control over participating advertisers and brands
- Custom deal packaging, allowing inventory to be bundled by format, audience, or content category

PMPs offer a middle ground between open RTB and programmatic direct, balancing scale with quality control.

A Comparison Table of the Above Media-Buying Processes

Below is a comparison table that illustrates how the three ways of purchasing media stack up against each other.

Media Execution Type	Programmatic			Non-Programmatic (Direct Campaigns)
	Real-Time Bidding (RTB)		Programmatic Direct	
	Public Auctions	Private Marketplaces (PMP/Deal ID)		
Price	Auciton	Auction and / or Deal ID terms	Pre-defined	Pre-defined
Direct Advertiser - Publisher relationship	No	Yes	Yes / limited	Yes
Inventory Volume	Non-guaranteed	Non-guaranteed	Guaranteed	Guaranteed
Inventory	All inventory that the publisher decides to put on public auction	Premium inventory	All, including premium inventory Bulk inventory (sometimes robust targeting is available)	All, including premium inventory Bulk inventory with limited targeting
Delivery	DSP / over RTB pipes	DSP / over RTB with Deal ID set	Programmatic-direct platform integrated with the publisher's ad server	Email / phone, manual ad tags entered in the publisher's ad server
Advantages	• Per-impression buying process • Advertisers and publishers use a single dashboard • Easy testing and adjusting Insights • Ability to sell remnant ad space	• Transparency of purchased inventory and pricing • Programmatic efficiency without middlemen • Becoming an industry standard • Can remove the need for direct-sales team	• Transparency • Automation • Better insights and control • Directness • Garanteed inventory volume	• Transparency • Direct advertiser-publisher relationship • Insights and better control • Garanteed inventory volume
Disadvantages	Limited access to the premium inventory	More expensive inventory	Direct deals and pre-defined prices may lead to overpaying for inventory	• Not easily scalable • Slow process

The Publisher's Waterfall

Waterfalling, also known as a daisy chain or waterfall tags, is a process used by a publisher to sell all remnant inventory. This process occurs when a publisher has been unable to sell its premium ad slots that are usually reserved for direct ad sales between the publisher's internal sales team and advertisers.

The term *waterfalling* comes from its tiered setup: each demand source (such as an ad network or SSP) is called in sequence, one after another, until the impression is sold. If the first buyer declines or fails to meet the price floor, the next in line gets the opportunity — creating a cascading "waterfall" effect.

The Publisher's Dilemma: High CPM or High Fill Rate?

Although ad networks help publishers monetize their remnant inventory, they still face what's known as *the publisher's dilemma*: should a publisher sell inventory at a high CPM and risk leaving some ad slots unfilled — losing potential revenue?

Or should they aim for full inventory fill at lower CPMs — sacrificing higher potential earnings for guaranteed sales?

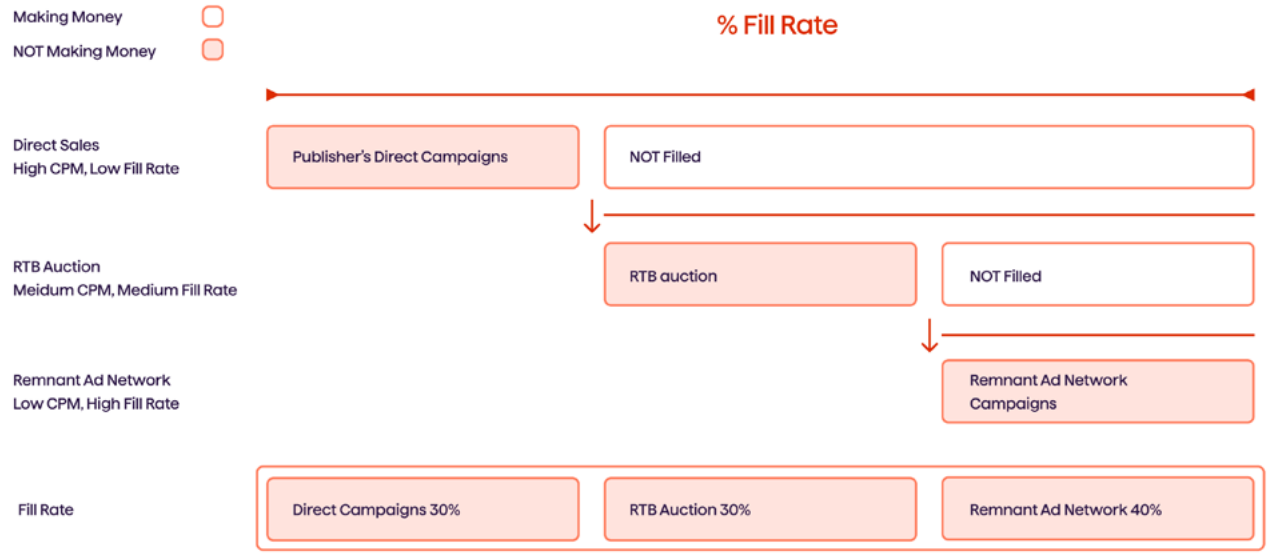

In the image above, we can see the publisher first tries to sell its inventory via direct sales, as these generally offer the highest cost-per mille (CPM).

If those slots remain unsold, the impression flows down the chain — first to SSPs and ad exchanges, where it is offered via open RTB auctions.
If the ad space is still not sold, it is finally offered to ad networks as lower-tier demand sources.

How Is Waterfalling Implemented?

The AdOps team sets up an ad network with a tracking tag that will execute if and when an impression is not filled.

This is typically configured in the publisher's ad server as well as in each ad network's system that the publisher works with in a field called Fallback Ads, Passbacks, Redirects, Default Ads, or similar.
Each passback must be configured separately for every ad network the publisher partners with. In practice, this means the AdOps team would instruct the premium ad network to "pass back" any unsold impressions to the remnant ad network, ensuring that no inventory goes unmonetized.

How Does Waterfalling Work?

If a publisher is unable to sell its direct buys, its ad server executes the first ad network's tag. There are a couple of possible outcomes, so let's take a look at a few likely scenarios.

Scenario 1

A direct deal with advertiser #1 can offer an impression for this ad call, so it sends an ad back to the user's browser.

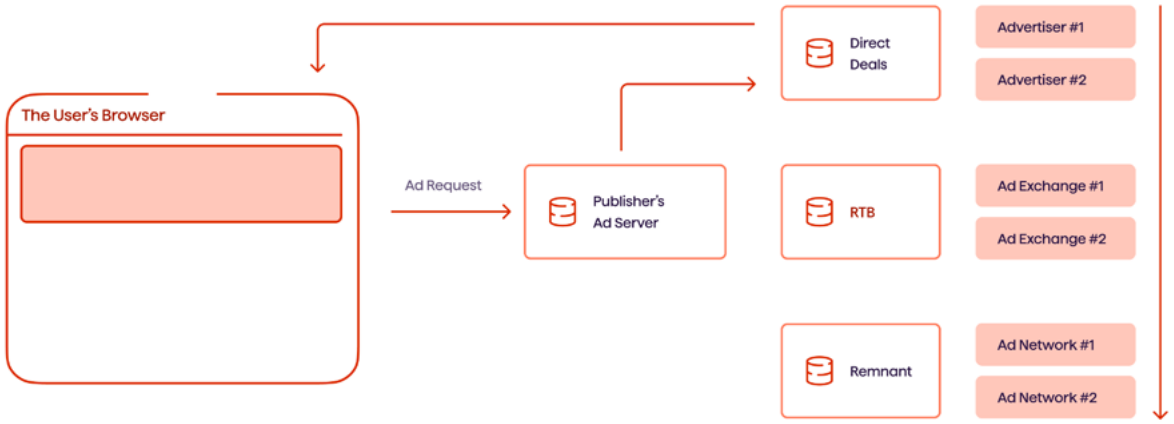

Scenario 2

The publisher's direct deals have no available impression to serve, so the system triggers tags to initiate RTB auctions through ad exchanges #1 and #2. When neither exchange returns a winning bid, the process continues by calling the remnant ad network #1, which successfully serves an ad for this request.

The winning creative is then rendered in the user's browser, completing the ad delivery process.

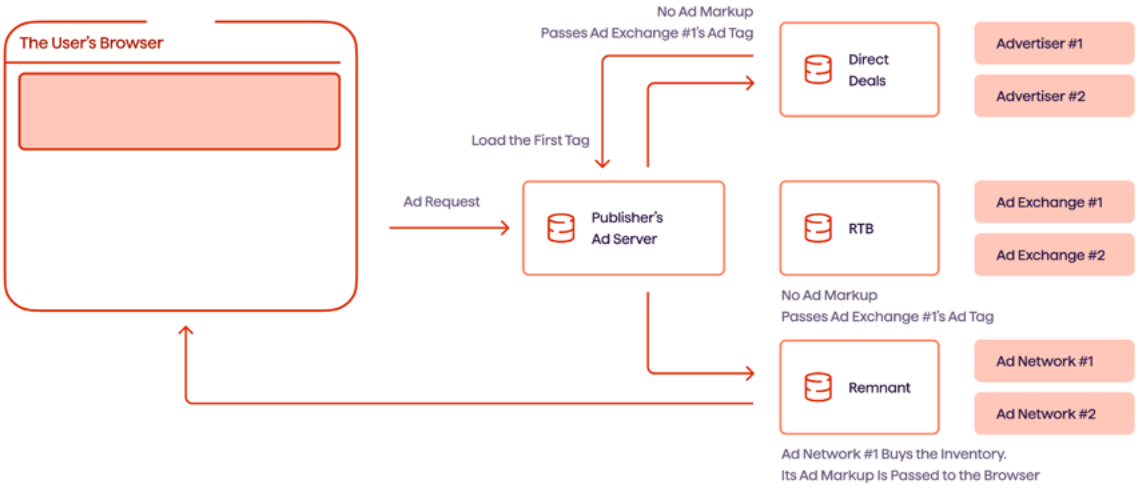

Why didn't any of the direct deals or RTB sources buy the inventory?

There are several possible reasons why these demand sources failed to return an impression:

- **No matching campaigns** — The advertiser's targeting criteria may not have aligned with the website's content or the user's profile, meaning the demand source (DSP or ad network) had no suitable ad to serve.
- **High floor price** — The publisher's set floor price might have exceeded the advertiser's bid threshold, resulting in no bids being placed.
- **Impression capping** — The advertiser may have already reached its frequency limit for that user within the specified timeframe, preventing additional ad delivery.
- **Timeouts** — The advertiser's ad server, DSP, or network could have taken too long to respond to the ad request. In such cases, the system times out, fails to return a passback tag, and the waterfall chain ends without displaying an ad — one of the key drawbacks of the waterfall method.

Scenario 3

If all the demand sources don't offer an impression, the publisher activates its fallback option. In most cases, this is an ad promoting its own products or services.

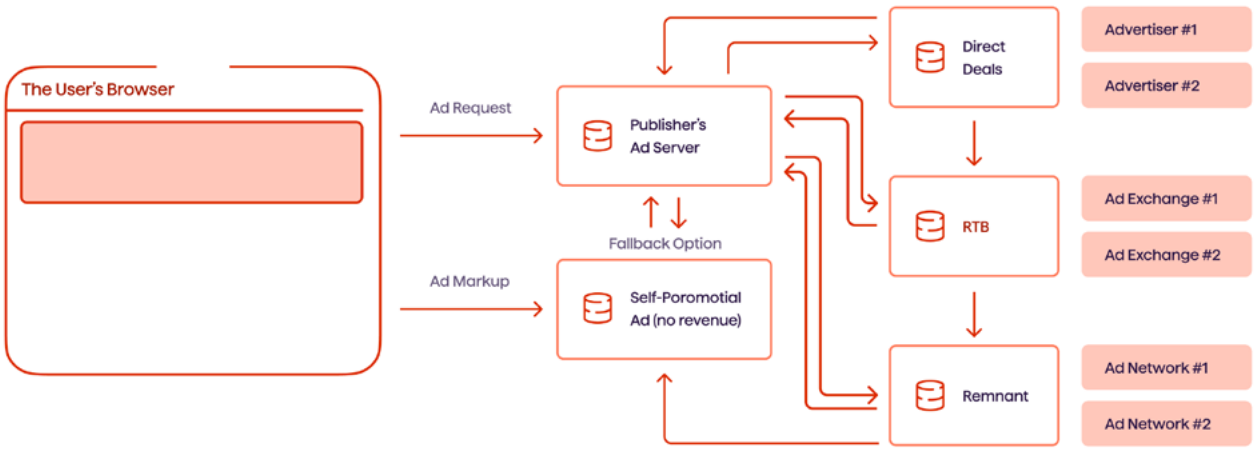

How Does the Ad Server Know Which Ad Network to Load First?

Publishers typically configure a ranking system that prioritizes demand sources based on their average historical yield — the average revenue each source has previously generated for the publisher. The ad server then calls these demand sources sequentially, starting with the one that has delivered the highest historical earnings.

However, this static hierarchy has a major limitation: it doesn't account for real-time fluctuations in bid values. A demand source further down the chain might be willing to pay significantly more for a particular impression than its historical average suggests.

For instance, a remnant ad network ranked third in the waterfall might have an average historical yield of $2 CPM, yet for a user who matches its ideal targeting profile, it could be willing to bid $5 CPM. Because of the waterfall's sequential logic, that higher bid may never be seen.

This inability of waterfalling to reflect real-time market value for impressions led to the emergence of header bidding — a more dynamic, simultaneous auction process that allows all demand partners to compete at once.

Header Bidding

Header bidding — also known as pre-bidding, advance bidding, or holistic yield management — is a programmatic technique that allows publishers to simultaneously solicit bids from multiple demand sources (e.g., DSPs, ad exchanges, or networks) before their ad server calls other tags such as direct deals.

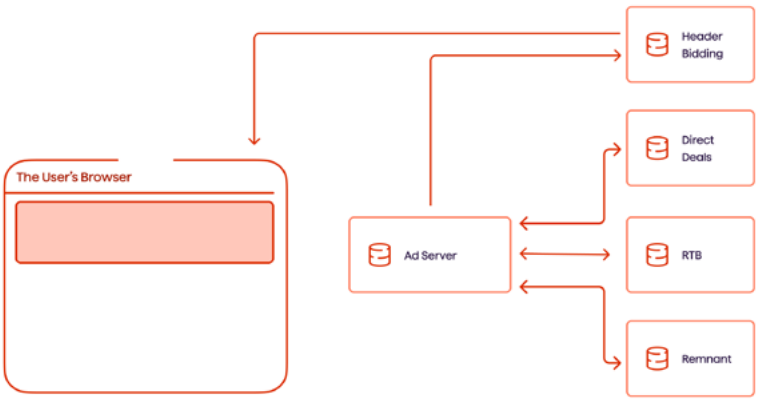

The process is executed through a JavaScript snippet placed in the website's header section — hence the name *header bidding*. When a user loads the page, this script triggers real-time auctions among participating demand partners, collecting bids within milliseconds.

Header bidding emerged to solve two major issues:

1. The inefficiency of waterfalling, where higher bids could be missed due to sequential auction logic.
2. Google's preferential treatment of its own ad exchange within its ad server (formerly DoubleClick for Publishers, now Google Ad Manager), which gave AdX an advantage over competing exchanges.

As a result, non-Google demand sources were often excluded from premium inventory, even when they were willing to pay higher CPMs. Header bidding leveled the playing field by allowing all demand partners to compete simultaneously for each impression.

How Does Header Bidding Work?

To implement header bidding, publishers need to add a piece of JavaScript code (aka snippet or tag) in between the <head></head> tags on their website.

This JS code often comes in the form of a wrapper (aka container), which is typically provided by SSPs and ad exchanges.

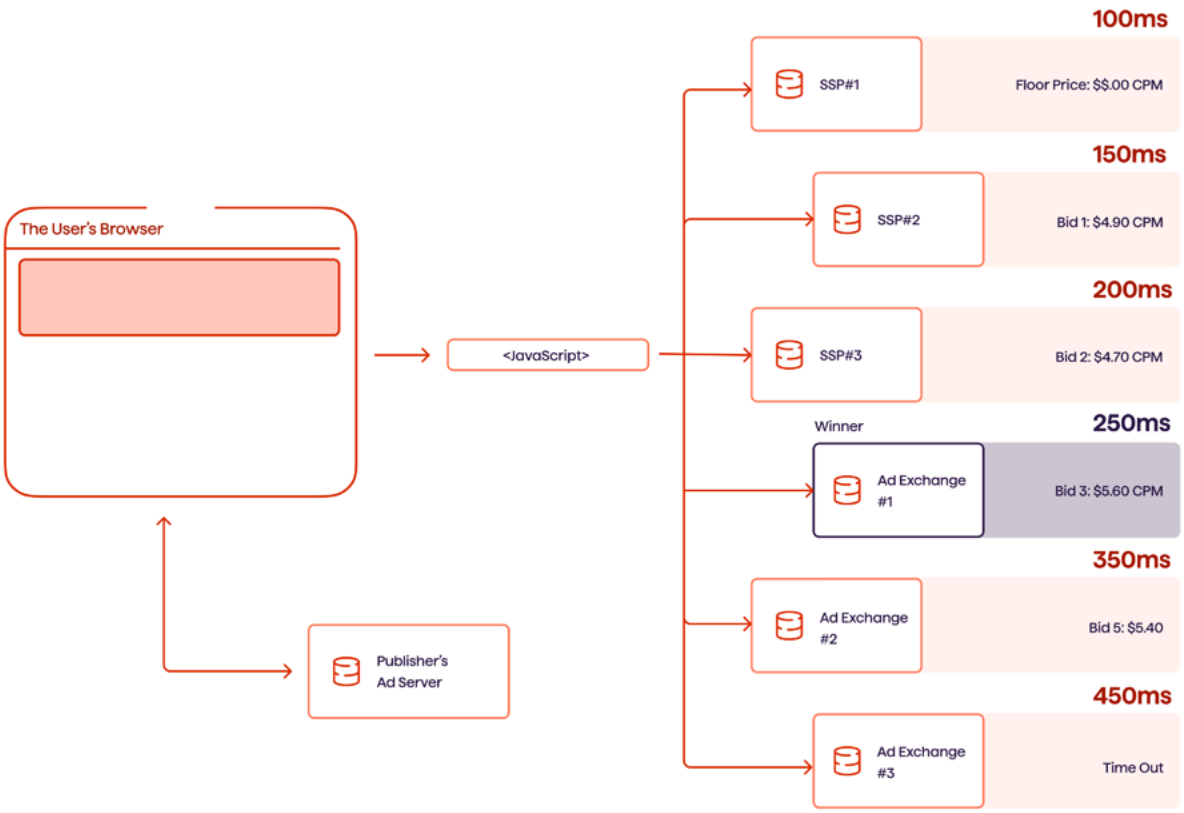

Here's what is happening in the image above:

- A user opens their web browser and types in the publisher's URL (e.g. publisher.com)
- The browser starts loading the page
- The header-bidding JavaScript code or wrapper located in the <head> tag executes and sends a request to the various AdTech platforms (SSPs and ad exchanges)
- The SSPs and ad exchanges send bid requests to multiple DSPs
- The DSPs analyze the bids and return a bid response if the impression matches their campaigns
- The highest bidder wins
- The bid passes on to the publisher's ad server and competes with other campaigns, such as direct deals
- If the DSP's bid is higher than the publisher's other campaigns, it is displayed to the user

Just like other media-buying processes, latency is a big issue with header bidding.

If a DSP, SSP, or ad exchange doesn't respond to the ad or bid request in time, they will be timed out and won't be able to submit a bid.

Timeout thresholds vary depending on the device and user connection speed. On desktop and laptop environments, the typical range is 400–800 milliseconds, while on mobile devices, where network conditions are less stable, the range extends to 800–1,200 milliseconds.

Because header bidding involves multiple simultaneous requests, optimizing timeout settings is crucial to balancing auction fairness, page-load speed, and revenue performance.

Prebid.js – Making Header Bidding Easier for Publishers

Prebid.js is a 100% free and open-source JavaScript framework designed to make it easier for publishers to run pre-bid auctions and get access to more demand with minimal integration hassle. Available at prebid.org.

How To Implement Header Bidding: Client-Side vs. Server-Side

When it comes to implementing header bidding, there are two primary approaches: client-side header bidding (CSHB) and server-side header bidding (SSHB).

- **Client-side** header bidding collects bids directly within the user's web browser (the "client")
- **Server-side** header bidding sends the bid requests to a dedicated server, which handles the auction externally

In both cases, the publisher must still embed a wrapper or JavaScript snippet on their website to initiate and coordinate the auctions.

Case Study

Header Bidding Control Center — We created an all-in-one solution for managing your Prebid.js setup without code changes. It allows you to change the configuration of ad units, select different demand partners, upgrade the Prebid.js version, and more.

Three main use cases for HBCC: client inventory management, changes to Prebid.js setup, Prebid.js analytics.

Below is a side-by-side comparison illustrating how client-side and server-side header bidding differ in their implementation and data flow:

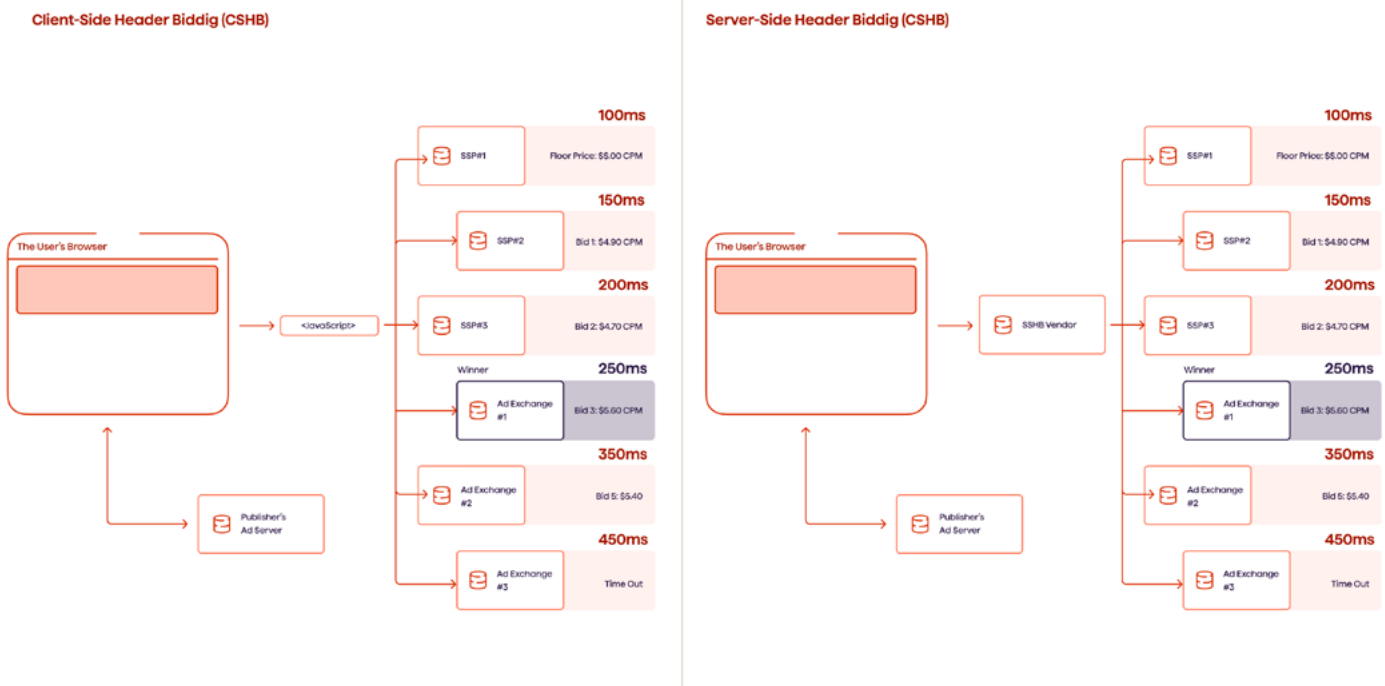

	Client-Side Header Bidding	Server-Side Header Bidding
Advantages	• Higher cookie-matching rates due to in-browser processing • Greater control over wrappers • More transparency into demand sources and clearing prices	• Reduced page-load latency with single server call • Ability to receive more bids due to fewer technical limits
Disadvantages	• Slower page load, affecting user experience • Potential browser compatibility issues • Limited number of browser requests	• Less control and transparency into pricing and demand • Harder cookie matching, reducing user addressability and possibly lowering revenue

Waterfall vs. Header Bidding: Benefits and Drawbacks

Below, we summarize the main benefits and drawbacks of both the waterfall method and header bidding for publishers that want to increase fill rates and maximize yield.

Benefits of the Waterfall

- Monetizes remnant inventory — Helps sell leftover ad space that might otherwise go unsold.
- Ease of implementation — Waterfalling is simpler to set up than header bidding and requires less technical expertise. Publishers only need to configure ad network tags within their ad server.

Benefits of Header Bidding

- Higher competition and pricing — Publishers receive bids from a broader range of buyers, some of whom may be willing to pay more than those connected directly to the ad server.
- Improved fill rates — Increases the likelihood of filling all available inventory — both premium and remnant — due to more participating demand sources.
- Greater market transparency — Publishers gain insight into true inventory value. For example, if a publisher's floor price is $1.50 CPM but header bidding yields an average of $2.00 CPM, it reveals the real market worth of that inventory.

Drawbacks of the Waterfall

- Lower yield potential — Demand sources are prioritized based on historical averages rather than real-time market prices, which can reduce revenue opportunities.
- Latency issues — Each tier in the waterfall adds delay, lowering the chances of the user actually seeing an ad.
- Revenue loss from timeouts — Some demand sources may fail to load passbacks or time out, leading to unfilled impressions.
- Complex setup management — Each demand source may require separate passback configuration, making ongoing maintenance time-consuming for AdOps teams.

Drawbacks of Header Bidding

- Latency from added scripts — Although it reduces passbacks compared to waterfalling, header bidding can still slow page-load times, especially in client-side implementations.
- Browser compatibility challenges — Client-side header bidding must remain backward-compatible across browsers, increasing maintenance complexity.
- Duplicate bid requests — Using multiple header partners can cause the same impression to be offered multiple times, wasting processing resources.
- Performance impact — Additional logic can slow down page rendering — usually negligible on modern devices but noticeable on older hardware and smartphones.

Auction Dynamics: First- and Second-Price Auctions and Hard and Soft Floor Prices

Auctions are a foundational component of real-time bidding (RTB) in programmatic advertising. Just like auctions for houses or artwork, online ad inventory is bought and sold to the highest bidder, but in milliseconds.

Understanding how different auction types and pricing strategies work helps both advertisers and publishers optimize their performance and revenue.

Second-Price Auctions (2PA)

In a second-price auction — also known as a Vickrey auction — advertisers place bids for a single ad impression. The highest bidder wins, but instead of paying their full bid, they pay just $0.01 more than the second-highest bid. This final amount is known as the clearing price.

For example:

- Bidder A offers $5.00
- Bidder B offers $4.00
  → Bidder A wins but pays $4.01 instead of $5.00.

The difference between the winning bid and the amount paid ($0.99) is referred to as the reduction or consumer surplus.

Second-price auctions became the standard model in RTB (real-time bidding) when programmatic advertising first emerged, offering predictability, fairness, and cost efficiency for advertisers.

First-Price Auctions (1PA)

In first-price auctions, the winning advertiser pays exactly what they bid on. If an advertiser bids $6.00 CPM and wins, the clearing price is $6.00 CPM.

This model gained popularity around 2017–2018 when many SSPs and ad exchanges transitioned from second- to first-price auctions. The shift aimed to increase transparency and counter the inefficiencies introduced by header bidding.

Bid Shading

The transition to first-price auctions led advertisers to pay more for impressions than they were used to under the second-price model. To offset this, many AdTech platforms introduced bid shading.

Bid shading uses algorithms to estimate the optimal bid price — close to what the advertiser would have paid in a second-price auction — by analyzing historical data, win rates, and market conditions.

While bid shading can help advertisers control costs, it introduces transparency challenges.

Some AdTech vendors may adjust shaded bids to increase their own margins, leaving advertisers unable to verify whether they're paying fair market value for impressions.

Moreover, because bid shading reduces the final bid amount, it can negatively impact publisher revenue.

Floor Prices

Floor prices are minimum CPMs that a publisher is willing to accept for their inventory. These prices are often defined in contracts between publishers and their AdTech partners to prevent inventory from being undervalued.

Hard Price Floor

A hard floor is a strict minimum. Any bid below this threshold is automatically rejected. For instance, if the hard floor is $4.25, bids of $4.24 or less won't be considered at all.

Soft Price Floor

A soft floor has a more flexible threshold. Bids between the soft and hard floor can still compete but may be treated under a different auction model. For example:

- Bids between the soft and hard floor may enter a first-price auction
- Bids above the soft floor can participate in a second-price auction

This approach allows publishers to capture value from borderline bids while maintaining pricing control.

8. User Identification

Accurate user identification is key to targeting, personalization, and measurement in AdTech.

This chapter explores how advertising platforms identify and recognize users across multiple sessions and devices — using tools like cookies, device identifiers, and universal IDs, while also addressing the technical limitations and privacy regulations that shape modern user tracking.

Cookie-Based Tracking

Cookies (aka web cookies, HTML cookies and browser cookies) are small files that are placed on a user's device by a web server when accessing websites.

Web cookies were created by Lou Montulli in 1994 as a way to remember stateful information in an otherwise stateless environment.

In simple terms, the HTTP protocol — which governs communication between web browsers and servers — is inherently stateless. It cannot retain information between sessions; it simply receives requests and sends responses.

Cookies solve this limitation by allowing browsers to store and send small pieces of data back to the server with each request, enabling persistent sessions, user preferences, and other personalized web experiences.

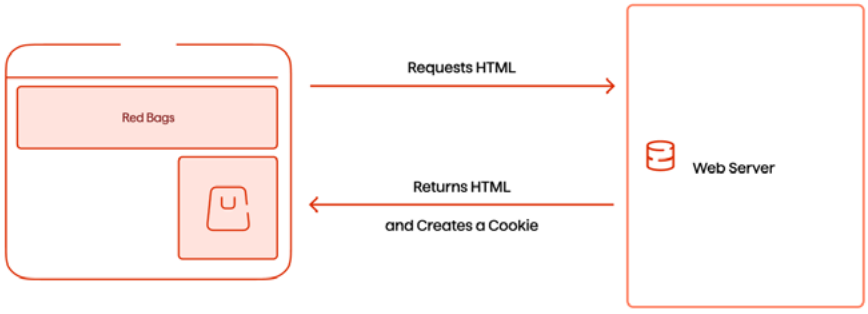

When a user revisits a website, cookies enable the site to recall key details from their previous session — such as which pages they accessed, what content they viewed, and even their personal preferences.

Some of the main uses of cookies include:

- Website personalization — Cookies store user preferences like language, theme, and currency, allowing a customized browsing experience.
- Authentication — To keep users logged in to their accounts, a unique session ID is stored in a cookie. This ensures users don't have to re-enter credentials each time they open the site.
- eCommerce functionality — Cookies used by eCommerce stores help web browsers remember which products users viewed, added to the shopping cart and purchased.
- Analytics and measurement — Cookies are used to store a user identifier that collects data about the user's interaction with the website under one profile and session, such as which pages they visited, what areas they clicked on and if they completed any goals (e.g. downloaded an eBook).
- Behavioral targeting and advertising — AdTech platforms use cookies to identify users and show relevant ads to them based on their previous behavior, such as which websites and pages they've visited. Cookies also help advertisers and publishers know which ads they've viewed and clicked on.

Case Study

We built a proprietary Consent Management Platform (CMP) over six years, releasing its core features, including a website cookie scanner, a cookie-category knowledge base, and IAB TCF integration. Today, the CMP handles several million consent requests daily, serving thousands of customers worldwide.

The CMP handles +10 million consent requests daily.

Cookies have been the most common method for identifying users on web browsers since the early days of the Internet. However, the introduction of privacy laws such as the European Union's General Data Protection Regulation (GDPR), and browser privacy features like Safari's Intelligent Tracking Prevention (ITP) have increasingly restricted cookie creation and access.

While all cookies function similarly, they fall into two main categories: first-party and third-party cookies. The distinction lies in the relationship between the website the user visits and the server that issues the cookie.

Third-Party Cookies

Third-party cookies — also known as tracking cookies or third-party trackers — are created by domains other than the one a user is currently visiting.

For example, if you visit *techcrunch.com* and the page loads a JavaScript file from an AdTech platform such as ssp1.com, a first-party cookie would be created for *techcrunch.com*, while a third-party cookie would be created for *ssp1.com*. Because *ssp1.com* is not the domain the user is visiting, its cookie is classified as a third-party cookie. Any other cookies created by external domains would also fall into this category.

For years, third-party cookies have served as the backbone of online advertising, enabling AdTech companies to identify users across different sites, deliver behavior-based targeted ads, and attribute impressions, clicks, and conversions. However, their effectiveness is declining due to stricter privacy laws like the GDPR, browser-level privacy features such as Safari's Intelligent Tracking Prevention and Firefox's Enhanced Tracking Protection, as well as privacy-focused browser extensions and plugins.

First-Party Cookies

First-party cookies are created by the domain (website) a user visit directly. For example, if you visit techcrunch.com, then it will create some first-party cookies and save them to your device.

These cookies are primarily used to enhance the user experience by remembering preferences and session details. For instance, first-party cookies help websites recall a user's selected language, keep items in a shopping cart, and maintain login sessions, so users don't have to re-enter their credentials each time they visit.

While first-party cookies can support some forms of online advertising, they are limited in their ability to recognize users across different domains. A cookie created by *ssp1.com* on *techcrunch.com* cannot be accessed by *ssp1.com* on another site.

To overcome this limitation, AdTech companies use third-party cookies and a process known as cookie syncing, which maps one user ID (stored in a cookie) from one platform to another (for example, from a DMP to a DSP), enabling cross-domain identification.

Modern Identity Solutions

There are several reasons why online advertising needs an alternative ID solution to replace cookies:

- Cross-site identification is essential — Being able to recognize users as they move across websites is critical for targeting, measurement, and attribution throughout the advertising ecosystem.
- Browsers lack a persistent user ID — Modern web browsers don't generate or expose a consistent identifier that can be used across sessions or sites.
- Third-party cookies are disappearing — Privacy regulations and browser restrictions are phasing out third-party cookies, reducing advertisers' ability to track users reliably.
- Cross-domain visibility is limited — Cookies can only be read by the domain that created them, preventing seamless tracking across multiple publishers.
- Cookie syncing is inefficient — Exchanging user IDs between AdTech platforms consumes resources, adds latency, and creates synchronization errors.
- Syncing isn't always accurate — Mismatched or missing cookies can result in incomplete user profiles and wasted impressions.
- Walled gardens have an advantage — Platforms like Google and Meta rely on deterministic, login-based identifiers ("people-based IDs"), giving them a major data advantage.
- Mobile tracking faces new barriers — Access to device identifiers (e.g., Apple's IDFA) has become more restricted, especially on iOS, complicating cross-device targeting.

Universal IDs

Universal IDs are unique identifiers that enable AdTech companies to recognize users across multiple websites and devices. These IDs are interoperable across platforms like SSPs and DSPs, allowing consistent user identification throughout the advertising ecosystem.

They serve the same purpose as third-party cookies — tracking, targeting, and measurement — but differ in how they're created. Universal IDs are derived from first-party data sources such as cookies and device identifiers (e.g., mobile IDs from smartphones or tablets), making them more privacy-compliant and resilient to browser restrictions.

Some universal IDs operate within one environment, such as web browsers, while others aim to identify users across different environments, such web browsers and mobile devices. For the latter, device graphs are used to match together the IDs generated in web browsers with the ones generated in other devices, e.g. mobile IDs in smartphones.

Privacy Sandbox

In response to the decline of third-party cookies, Google introduced Privacy Sandbox — a set of standards to support advertising use cases while minimizing user-level tracking.

The initiative aims to move digital advertising into a more privacy-conscious, controlled environment — one that balances user protection with the continued viability of an ad-supported web. It addresses the growing concerns of users, privacy advocates, and regulators over how personal data is collected and shared.

At the center of this framework is the Topics API, which enables interest-based advertising. This approach allows ads to be selected based on a user's inferred interests — determined by their recent browsing activity or, on Android, their app usage.

Unlike contextual advertising, which targets ads based on the content currently being viewed, interest-based advertising focuses on the user's broader behavior and preferences.

This method benefits both advertisers and publishers:

- Advertisers can better reach users who are likely to be interested in their products or services
- Publishers can combine interest data with contextual information to more effectively monetize their content through advertising
- Even platforms with non-commercial or hard-to-monetize content can generate revenue by serving relevant ads based on users' interests

The Topics API introduced a new approach to interest-based advertising. Instead of tracking individuals, it assigns users general interest categories (like "Fitness" or "Technology") based on their recent browsing behavior.

When the user visits a participating site, the browser shares a small, rotating selection of these topics with advertisers to help serve relevant ads — while avoiding the need to share personal identifiers or detailed browsing history.

Clean Rooms

A data clean room is a piece of software that allows brands and advertisers to run targeted advertising campaigns, apply frequency capping, measure and report on the performance of campaigns, and run attribution — all in a privacy-friendly way.

This is achieved by allowing companies to upload their first-party data and compare it against aggregated datasets from other participating organizations within the clean room.

Unlike other types of data partnerships whereby companies directly exchange user-level data, such as cookie IDs, device IDs, and IDs created from hashed email addresses, data clean rooms match the first-party data provided by brands and advertisers together but prevent any user-level data from being accessed outside of the data clean room. All first-party and user-level data stay within the data clean room and isn't shared with anyone else.

Both advertisers and publishers collect valuable first-party data from different sources. To help them collect and manage this data, they can use a custom data platform (CDP). A data clean room extends the capabilities of a CDP and takes data management to the next level.

A CDP enables the collection, sharing, and processing of first-party data, typically focusing on user-level identifiers and behavioral insights. A data clean room, on the other hand, shifts the focus toward the use of anonymized and aggregated data, ensuring that personally identifiable information (PII) remains protected.

Currently, the main alternatives to data clean rooms include Universal IDs, Google Chrome's Privacy Sandbox, and contextual ad targeting. Emerging solutions — such as crypto identities, which aim to represent individuals through encrypted avatars — also show potential for secure, privacy-preserving data matching and testing without sharing PII.

Examples of existing clean room solutions include Google Ads Data Hub, Amazon Marketing Cloud, and Snowflake Clean Room.

Mobile and App Identification

Mobile Device IDs (MAIDs)

Mobile IDs, also known as mobile advertising IDs (MAIDs), are unique identifiers associated with a user's mobile device — such as a smartphone or tablet. Because nearly every mobile device has one, these identifiers are more persistent than web cookies. While users can't disable or delete them entirely, they can easily reset them at any time.

MAIDs like Apple's Identifier for Advertisers (IDFA) and Google's Google Advertising ID (GAID) are widely used to track user behavior within mobile apps and support key advertising functions, including:

- Attribution — Tracking which advertisement led to an app install or in-app event
- Audience targeting — Building segments based on app usage and behavior
- Frequency capping — Ensuring users don't see the same ad too many times

However, recent platform-level privacy changes have significantly reshaped how MAIDs are used:

- Apple's App Tracking Transparency (ATT) framework requires explicit user consent before apps can access the IDFA. As a result, opt-in rates are low, reducing the effectiveness of IDFA-based targeting and measurement.
- Google's Privacy Sandbox for Android aims to phase out GAID, replacing it with more privacy-preserving solutions that limit cross-app tracking while maintaining essential ad functionality.

App-Specific Identifiers and SDKs

In addition to MAIDs, many apps use proprietary identifiers generated through Software Development Kits (SDKs). These IDs are often used internally or shared with trusted AdTech partners for audience development and analytics.

SDKs can provide valuable data such as session length, location, device type, and engagement metrics, but they must be implemented carefully to remain compliant with privacy policies and app store guidelines.

Cross-Device and Cross-Channel Identity Resolution

ID and Device Graphs

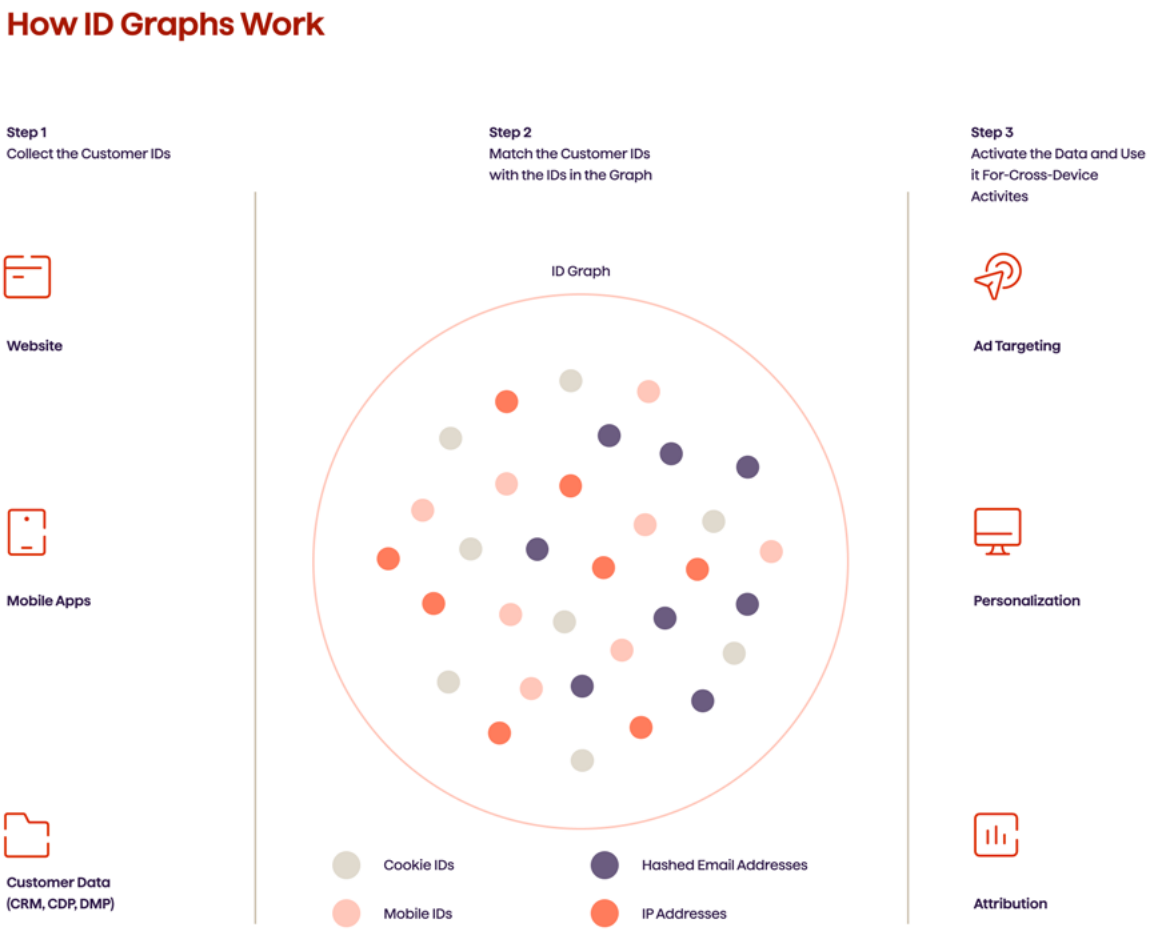

ID graphs map various identifiers (cookies, MAIDs, emails, IP addresses) to a unified user profile. This allows marketers to deliver consistent messaging across devices and platforms.

Deterministic and Probabilistic Matching

Deterministic and probabilistic matching are processes used to identify users across different devices. Companies will often use both to increase match rates.

Deterministic matching involves creating a profile of users consisting of different pieces of data about them. These profiles are then used to identify users on different devices by looking for a common identifier.

Common identifiers can include:

- Email address
- First and last name (if uncommon)
- Address
- Date of birth
- Phone numbers

It's important to note that this information is typically hashed when collected to protect user privacy and remove personally identifiable information (PII).

In online advertising and marketing, the most common method of deterministic matching is through email addresses — a unique, cross-platform identifier often found in multiple datasets.

Companies like Meta, Google, and LinkedIn can perform deterministic matching with exceptional precision because users must create accounts and sign in with an email address to access their services across devices.

The main advantage of deterministic matching is accuracy. It's much more accurate than probabilistic matching; the most deterministic matching rates are around 80-90%.

The main drawback, however, is that it lacks scale, as most companies don't collect this type of data, and email addresses aren't typically used for buying and selling online advertising.

To address the issue of scale, publishers are requiring users to create an account or subscribe using their email address to access certain content.

The two main ways they can do this are:

- By way of encouragement — Publishers can encourage visitors to provide an email address in exchange for more access and content
- By way of force — Publishers can gate their content and restrict access to it unless users subscribe or create an account

These tactics would work well with large publishers like news sites, as they typically have an engaged audience that regularly visits their sites compared to small- and medium-sized publishers, as not everyone will want to create an account just to read a few blog posts.

Unlike deterministic matching that uses a common identifier, such as an email, to match users to devices and applications, **probabilistic matching** uses various pieces of data, algorithms, and statistical modeling to make a match.

The type of data used for probabilistic matching includes:

- IP address
- Location
- Interests, behavior, and browsing history
- Wi-fi networks

Although probabilistic matching is less precise than deterministic matching, it often leverages deterministic datasets to train algorithms and improve accuracy. This process involves exposing algorithms to a sample of deterministic and probabilistic data (hundreds of thousands of records), so they learn to identify matching users. Once trained, the algorithms can be applied to millions of records without deterministic identifiers.

While probabilistic matching cannot match the accuracy of deterministic methods, its main advantage is greater scale and reach.

However, it also has several drawbacks, including:

- Lack of transparency in matching methods and accuracy
- Redundant or outdated data due to limited oversight
- Reduced data availability caused by privacy regulations such as the GDPR, which requires user consent to collect identifiers like IP addresses and location data

What Are Deterministic and Probabilistic Matching Used For?

Because deterministic and probabilistic matching aim to identify users across different devices and apps, their main use cases are cross-device targeting and cross-device attribution.

Cross-device targeting involves recognizing users across different devices and serving ads based on their collective behavior — such as the websites they visit and the products they purchase. For instance, a user might browse a jacket on their laptop and later see the same jacket (or a similar product) advertised on their smartphone.

Cross-device attribution, on the other hand, focuses on connecting ad interactions with conversions that occur on different devices. For example, if a user clicks on an ad for running shoes on their smartphone but completes the purchase on their laptop, cross-device attribution would link the smartphone click to the laptop purchase.

Alternative Methods

Fingerprinting

Fingerprinting tracks users by analyzing browser and device attributes (e.g., screen size, fonts, plugins). Though effective in identifying users without cookies, it's controversial and restricted by many browsers due to its invasive nature.

Device fingerprinting is a technical process that aims to identify and track online users based on the characteristics of their devices. It works by gathering bits of information to create an identifier, which is then used to identify individuals across different websites.

While many different users may own the same device, each one will be configured slightly differently according to the user's individual preferences and requirements. Data about these configuration changes can be aggregated to create a recognizable "device fingerprint."

Information used to create a device fingerprint can include:

- Browser version
- Operating system
- Language
- Items installed (plugins, fonts, etc.)
- Location and time-zone settings
- Browser settings

Why Do Companies Use Device Fingerprinting?

Device fingerprinting emerged as a solution to several challenges AdTech and analytics companies faced regarding the availability and reliability of cookies.

Over the past decade, many users have increasingly deleted or blocked cookies through browser settings and ad-blocking extensions to protect their privacy. This shift has made it significantly more difficult for AdTech and analytics platforms to consistently identify and track users across the web.

While device fingerprinting is less precise than cookies, it serves as a valuable fallback method when cookies are blocked or deleted. It can also be used alongside cookies to improve the likelihood of accurately identifying users across different sessions and websites.

Contextual Targeting

Contextual targeting does not rely on user identity. Instead, it targets users based on the content they are consuming (e.g., an article about fitness).

As privacy regulations tighten and identity solutions become harder to scale, contextual targeting is enjoying a renaissance for being privacy-safe and effective.

Advantages	Disadvantages
It's more privacy-friendly compared to behavioral targeting (e.g., via IDs) as it allows advertisers to reach their audiences without needing to use user-level data	Advertisers can only show ads based on the context of the page, meaning they can't use any user-level data to power the targeting, which can cause them to show ads that aren't relevant to the user

9. Data In AdTech

Data is the backbone of AdTech, powering everything from targeting to measurement.

This chapter explains the three primary types of data — first-party, second-party, and third-party — including their sources, collection methods, and how they're combined to create comprehensive audience profiles. You'll also learn how AdTech platforms manage and organize this data to enable precise segmentation and more effective advertising.

First-Party Data

First-party data is considered the most valuable type of data for both advertisers and publishers because it is collected directly from people who have interacted with the brand, such as customers.

First-party data is typically collected through:

- E-commerce and offline transactions — Details about purchased products, order values, and personal information such as names, postal and billing addresses, email addresses, and phone numbers.
- Customer relationship management (CRM) systems — Data about people who have created an account, downloaded a digital asset (e.g., an eBook), or made a purchase. Like e-commerce data, this often includes names, addresses, phone numbers, and email addresses.
- Website and mobile app analytics — Tracks user behavior, such as pages viewed, videos watched, and other content interactions.

First-party data can originate from both online and offline sources. Brands and advertisers use this type of data to convert visitors into customers and upsell products and services to existing customers.

Second-Party Data

Second-party data is sometimes referred to as partner data, as it is first-party information collected by one company and sold or traded to another.

These partnerships usually involve two non-competing companies that cater to similar audiences.

For example, a hotel chain might partner with an airline and purchase or exchange the airline's first-party data. The hotel chain could then use that data to run targeted ad campaigns promoting its hotels to the airline's customers — a logical match since travelers booking flights are often also looking for accommodation.

The partnership can be one-way (the hotel chain simply buys the data) or reciprocal, with both companies exchanging information. In the latter case, the airline could also display ads and offers to the hotel chain's customers.

Because airlines collect highly valuable first-party data, they can form multiple partnerships — for instance, with luxury brands targeting their high-income travelers with exclusive products like watches or jewelry.

While first-party data remains more valuable because it includes known or engaged customers, second-party data help brands expand their reach to new, but relevant, audiences who share similar behaviors or interests.

Third-Party Data

In terms of value, third-party data ranks last among the three main data types. It's neither collected directly by the advertiser or publisher nor obtained through a data partnership. Even so, third-party data still adds value to marketing and advertising campaigns and provides a couple of advantages over first- and second-party data, with the ability to reach a much bigger audience being the main one.

Third-party data is typically supplied by data brokers or is added as a layer by a DMP vendor.

Many publishers and merchants monetize their traffic by adding third-party trackers to their websites or tracking SDKs to their apps. These tools collect information about audiences and send it to data brokers and DMPs.

The data collected may include:

- Browsing history and content interactions
- Purchase behavior
- Profile details voluntarily entered by users (e.g., gender or age)
- GPS geolocation and device data

From these datasets, data brokers infer additional information such as interests, purchasing intent, income group, and demographics.

Third-party datasets can also be further enriched with information from offline sources, like credit card companies, credit scoring agencies, and telecommunications providers.

The result is a collection of audience segments — aggregated user profiles built from multiple attributes such as interests, demographics, and location — ready for targeting and analysis across campaigns.

Where Is Data Obtained?

Brands, advertisers, marketers, and publishers gather data from a range of online and offline sources.

Online Sources	Offline Sources
• Analytics tools • Customer-relationship management (CRM) systems • Enterprise resource planning (ERP) systems • Marketing automation platforms • Mobile and web apps • Campaign analytics	• Point of sale (POS) • Offline CRM and ERP systems • Transactional data

The types of data listed above are typically stored across multiple databases — either within the advertiser's or marketer's own systems, or in the databases managed by their software vendors.

Combining Online and Offline Data Together

Companies that collect both online and offline data would combine them together to get a clearer picture of their customers and audience.

For large companies, such as retailers, integrating their offline and online records is not an easy task, but once it's done, it can prove valuable as it provides several business advantages.

If, however, a company collects small amounts of offline data (e.g. only email addresses), it is possible to just import the data into a database or DMP. But if a company collects large amounts of offline data, then they will need to onboard it into a data platform like a DMP or CDP.

The Data Fragmentation Problem

While collecting vast amounts of data from multiple sources allows advertisers to improve campaign performance, the data is often stored across multiple tools and platforms. These individual databases are often referred to as data silos (a collection of data controlled by one department and isolated from other departments within an organization).

When data is fragmented across silos, advertisers lose visibility into the full picture of their target audiences and campaign performance.

This fragmentation often leads to poor decision-making, missed opportunities, and unnecessary ad spending.

Case Study

Custom data management stack
— We developed a comprehensive data management stack featuring data collection and normalization, customer profile building, cookie syncing, low-latency APIs for retrieving visitor data, audience creation and export, and dynamic content formatting.

With this solution, the client gained real-time customer insights and the ability to deliver dynamic, personalized content.

Data Collection in a DMP

Collecting data can be done in a few different ways, depending on where the data is stored.

Pixels and Tags

Probably the simplest way for a DMP to collect first-party data is by adding a 1×1 transparent pixel (also known as a tag or tracking pixel) to your website.

The pixel itself is just a piece of HTML. When the pixel loads on a page, it sends a request to the DMP to retrieve the 1×1 transparent image.

Once the DMP returns the 1×1 pixel, it can assign a cookie to the user's browser and begin storing information associated with that cookie. The collected data is then passed to the DMP for processing and audience profiling.

Piggybacking

Piggybacking occurs when a single master pixel is placed on a site's pages to trigger multiple tracking pixels from various sources or networks that aren't directly embedded on the site.

When the master pixel loads, it automatically calls and loads the additional pixels, allowing multiple systems to collect data from a single tag execution.

Tags

Tags are pieces of JavaScript or iframe code. Similar to pixels, tags are added to a website and when loaded, send a request to a DMP. The DMP responds to the request and places a cookie in the user's browser and collects data.

Sometimes, publishers will use a tag management system (TMS) to control and manage the various JavaScript snippets and pixels they have on their website. These tags and pixels are placed in a container which is inserted into a website's pages, usually directly under the opening body element <body>.

The main benefit of a tag manager is that publishers can easily add, remove, and modify their HTML tags, JavaScript snippets, and pixels from a single user interface, rather than having to ask their web developers to make manual changes in the website's HTML.

Application Program Interface (API)

APIs are used to exchange data between web servers and DMP. This type of data exchange is ideal for companies that have a number of data silos as it allows them to efficiently collect data from different databases.

This form of data collection is also referred to as server-to-server integration.

First-Party Data Onboarding

First-party data onboarding is the process of integrating a company's offline customer data with its online customer data to create a unified customer profile.

So, for example, a company could have the following customer data in their offline database:

- Names
- Residential addresses
- Phone numbers
- Email addresses
- Dates of birth

And all other data they have about customers in their offline customer relationship management (CRM) and transactional systems.

They could then onboard it with the data they have in their online databases, such as data from their web-analytics tools and ad servers, account information (e.g. account information from the company's online payment system), or any other online information the company has collected about the customer, including the same information they've collected offline — e.g., name, email, and residential address.

Depending on the amount of offline data a company has, they might just be able to import the data as a CSV file. But if they have a large amount of data, which is often the case, then they would likely need to use a data onboarding platform like LiveRamp.

The general process involves a company uploading offline data with an onboarding platform, anonymizing it to remove any personally identifiable information (PII), e.g., email addresses, names, physical addresses, etc., and matching the offline data with the company's online data.

Data Normalization and Enrichment in a DMP

Once the data has been collected, it's time to normalize it.

The data-normalization process can include a number of the following actions:

- Gathering IDs from web cookies
- Deleting redundant or useless data
- Transforming the source's data schema to the DMP's data schema
- Enriching the data with additional data points, such as geolocation and OS/browser attributes

The data normalization and enrichment stage provide two main benefits:

- It organizes the various data sets into a common format
- It improves data value and quality

During the normalization and enrichment stage, each user will be assigned a unique ID and given different attributes, which will play a key role in the segmentation stage.

These attributes can include:

- Age
- Gender
- Location
- Browser history
- Interests
- Purchase history

Profile Building and Merging

A profile is a structured collection of data gathered from events tracked by a DMP. It represents a unique user and may contain the following attributes:

- Profile ID
- Cookie ID (list)
- Hashed email (list)
- sID / uuID (list)
- Country (last seen)
- Name (nullable)
- device_type (last seen)
- device_vendor (last seen)
- device_os (last seen)
- browser_vendor (last seen)
- Gender (nullable)
- Company (nullable)
- Company size (nullable)
- Matching IDs (list)

In some cases, a profile will be created containing only a few pieces of data (e.g., cookie id, device_type, and device_os) and will be extended when more data becomes available — i.e. **profile building**.

When a DMP receives new events containing identifiers already stored in the system, the associated profile is updated with the new data. Conversely, if an event introduces an identifier not yet recognized by the DMP, a new profile is created.

It's common for two or more profiles to share overlapping identifiers (for example, the same cookie ID). When this happens, the DMP performs a process known as **profile merging**, combining the overlapping profiles into a single, unified user record.

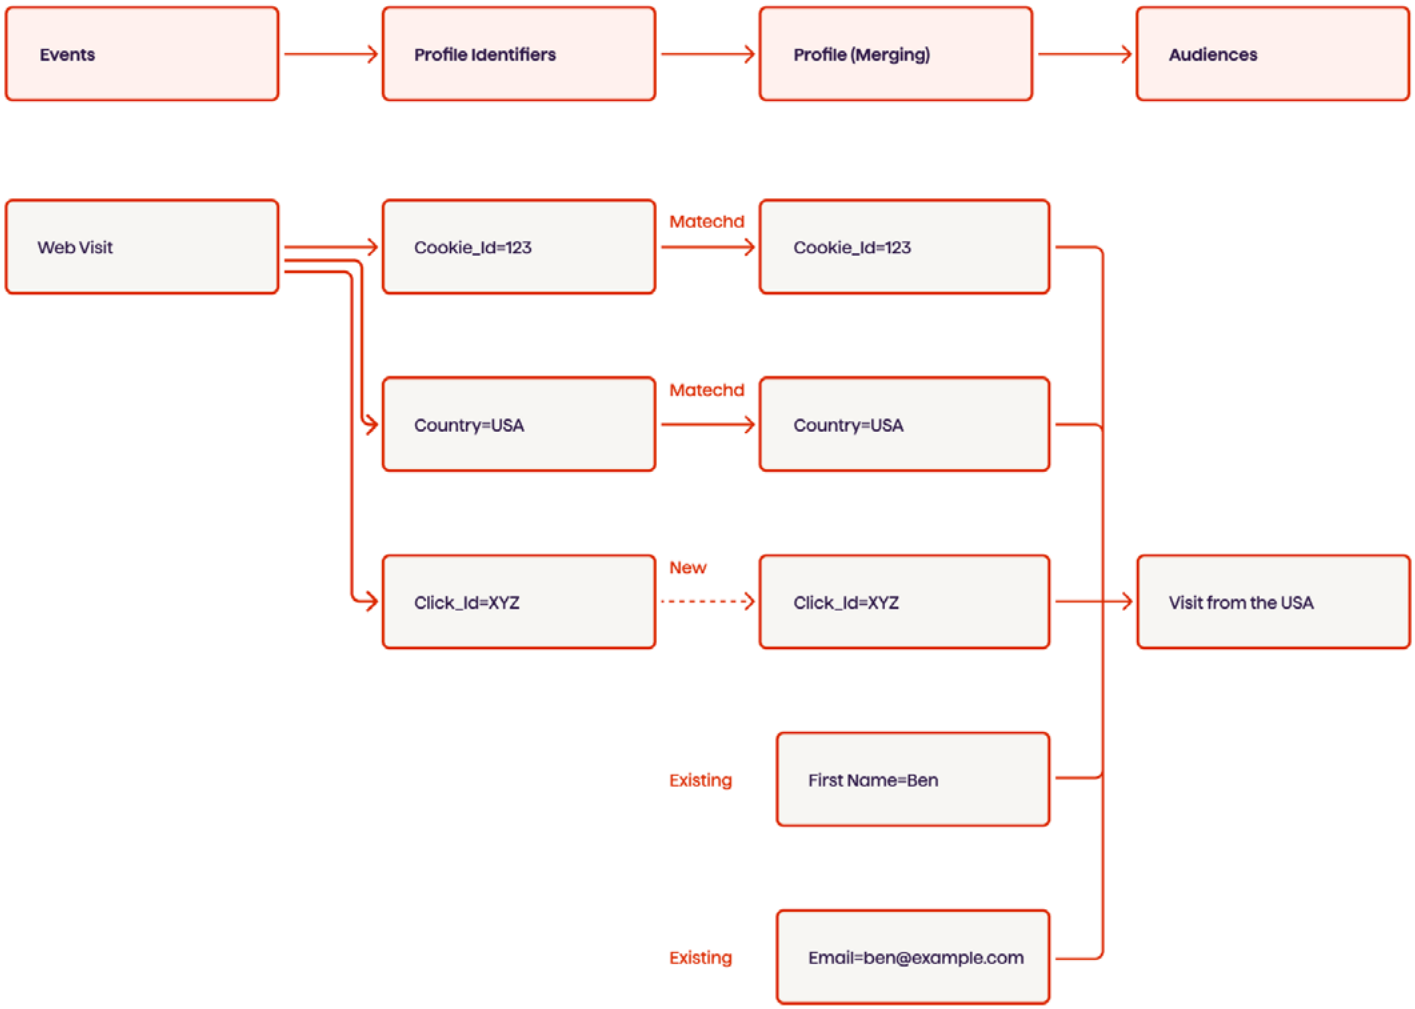

The goal of profile merging is to eliminate duplicates and ensure that no two profiles share the same unique identifiers (e.g., cookie IDs or email addresses).

Most DMPs would use a master ID — a single ID associated with one profile — to ensure accurate profile merging. Most often, this would be a persistent ID, such as an email address.

When new events containing the master ID enter the DMP, all other data associated with the event will be added to that profile.

Data Storage

Although the concept of storing data in a DMP seems rather simple, the actual technical implementation can be challenging due to the large amount of data being stored, as well as the need to move it to other areas and prevent data loss.

Data Taxonomies

Taxonomy in a DMP refers to the naming convention used for various pieces of data.

For example, instead of having two taxonomies like “user” and “visitor”, you could create or define one taxonomy (e.g. “user”) that represents both terms.

Below is an example of how an ecommerce store could structure its taxonomies:

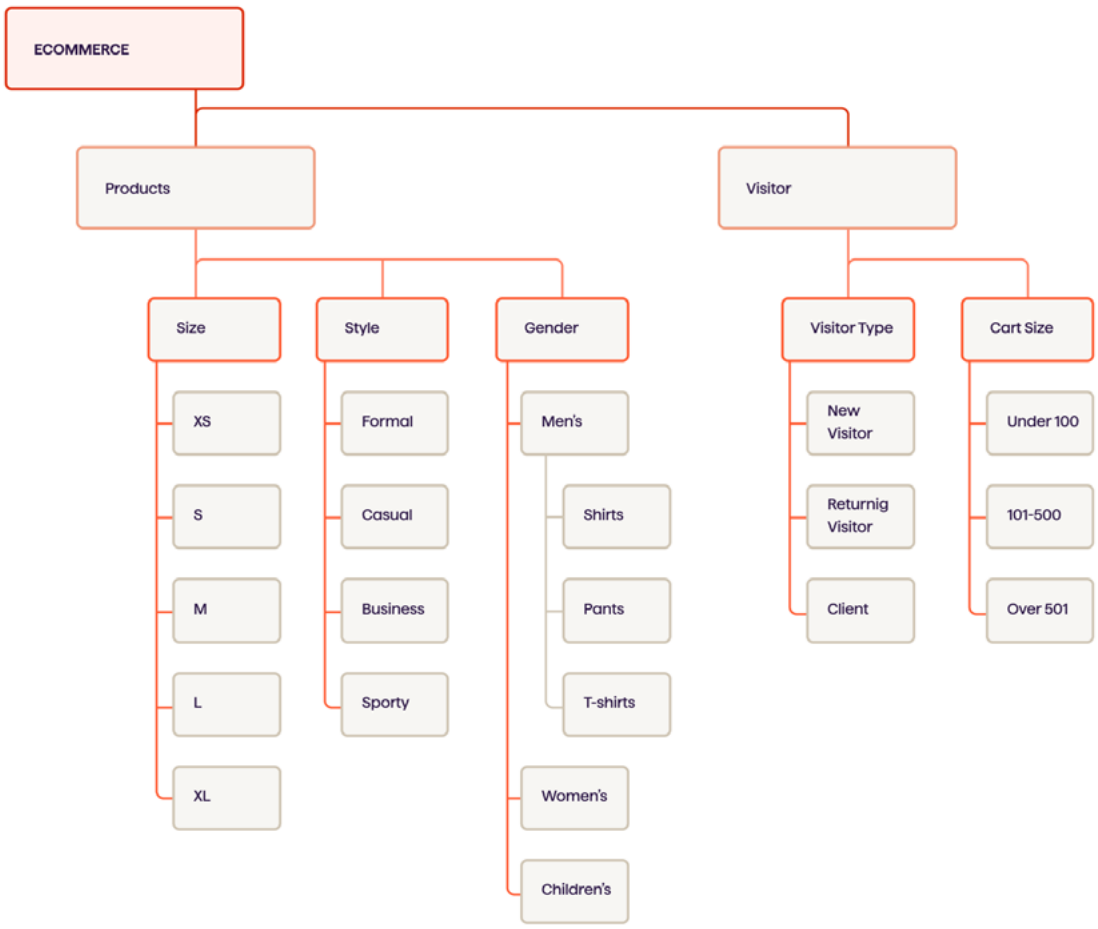

Audience Segmentation and Creation

Audience segmentation is the process of grouping users based on shared characteristics such as age, location, behavior, and interests. These segments form the foundation for data activation, enabling advertisers and marketers to use them for ad targeting, personalization, and analytics.

How Does a DMP Create Audience Segments?

To create audience segments, a DMP uses a series of conditions to filter the data and produce specific groups of users.

The conditions may include general information such as:

- Country, region, or city
- Device type
- Operating system
- Referral URL

As well as behavioral signals such as:

- Events (e.g., button clicks, page views)
- Conversions (e.g., downloads, purchases)
- Ads viewed

In addition, segmentation can incorporate demographic, for example:

- Relationship status: in a relationship
- Interests: gardening
- Age group: 35-39
- Gender: male
- Home value: between $200k – $400k
- Annual income: between $60k – $90k

Advertisers can then combine multiple segments to directly target the audiences they want to reach with their online advertising campaigns.

Here's an example of what that might look like:

Apart from selecting which users to include in an audience segment, you can also add filters to exclude users from the segment and set the recency and frequency of certain actions.

For example, you could add users who have viewed your website at least 5 times (frequency) in the past 30 days (recency).

These two additional factors will go a long way towards defining your audience's segments and their usefulness and can significantly impact the relevance and scope of the segments.

- **Relevance** — Increasing the time frame for data points to be included. For instance, setting a time frame of "greater than 30 days" can add users who may be less likely to convert. However, raising the frequency to "at least 3 times" for certain event information can mean adding a user who is highly engaged and likely to convert.

- **Scope** — Similarly, extending the time frame and reducing the frequency will broaden the scope of the audience, which would be useful for brand awareness but not for increasing conversions.

Once you've created audience segments, you can now activate your data.

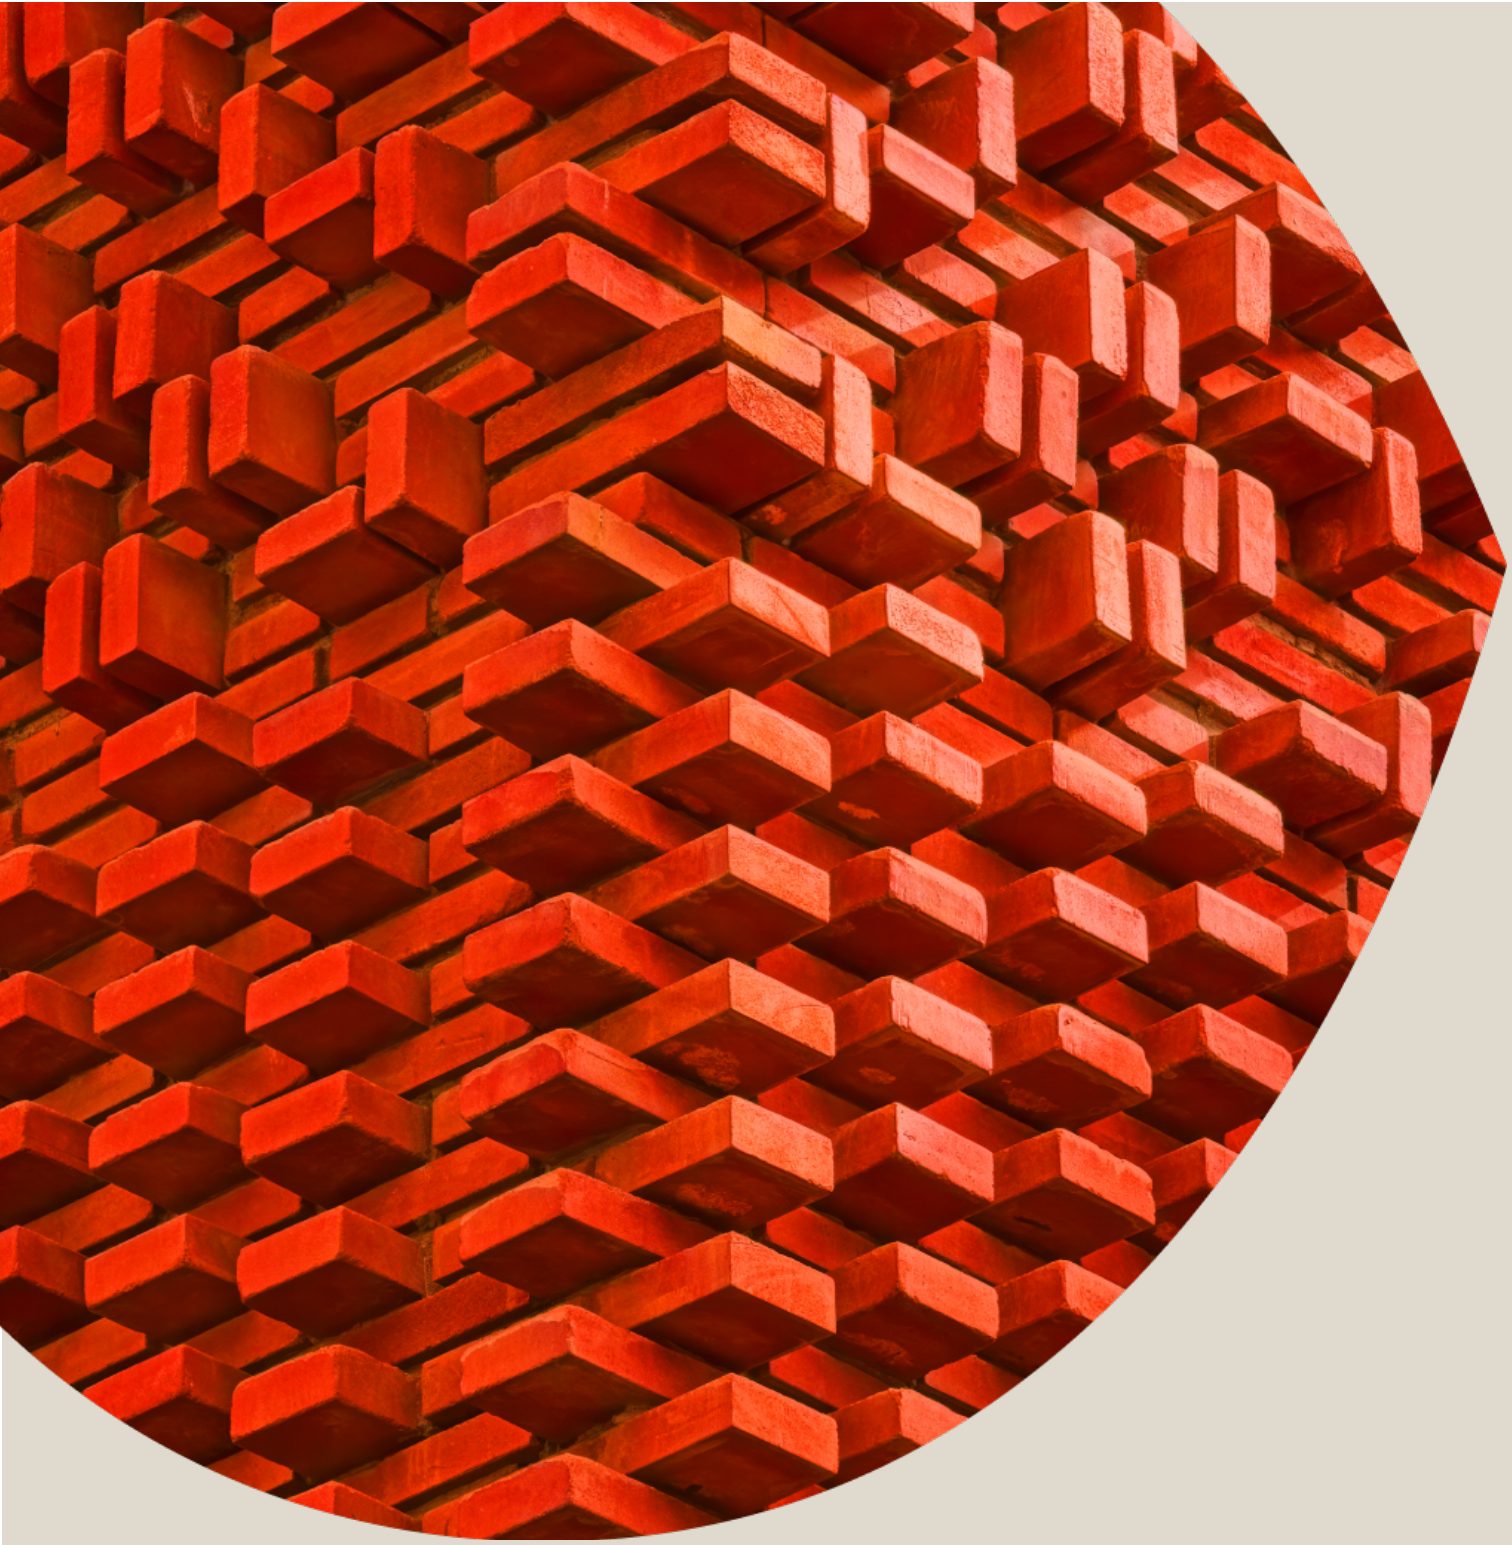

10. Attribution

Attribution is the key to understanding which marketing touchpoints influence a customer's journey before they convert.

This chapter explores how advertisers identify, track, and assign credit to those interactions across channels, devices, and even offline-to-online paths. You'll learn about common attribution models, cross-device tracking, and practical methods for connecting offline and online campaigns.

What Is Attribution?

Attribution is the process of identifying which touchpoints a consumer interacted with or was exposed to during a period before they completed a goal set by an advertiser or marketer. Attribution allows advertisers and marketers to make improvements to their campaigns by understanding which touchpoints are working and which ones aren't.

Attribution has always been a part of advertising and marketing, but it's through the use of data and technology that advertisers and marketers of today are able to more accurately attribute conversions.

Case Study

TV Ad Analytics Platform — We created a solution to measure TV ad impact on web traffic and conversions. The client's media agency gained next-day insights and handled hundreds of thousands of visits daily.

There were 2 primary components to creating this solution: a cookie-less tracking technology and web analytics platform.

Online to Online Attribution Models

Online to online attribution identifies which touchpoints a user had before they completed a goal across different online channels.

Since most online ad campaigns aim to drive users to a website, advertisers and marketers rely on attribution reports from web analytics tools, MarTech platforms, and AdTech platforms.

There are two main types of online-to-online attribution — inter-channel and intra-channel.

Inter-channel attribution looks at touchpoints across different channels.

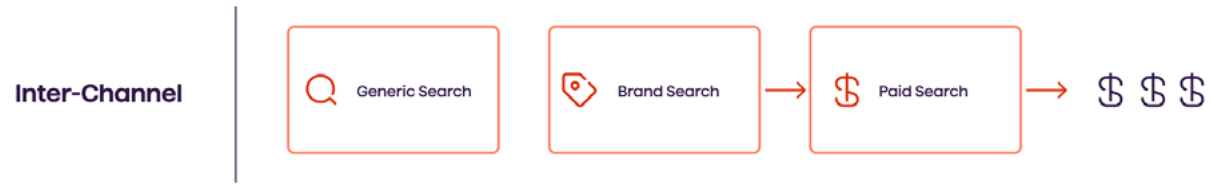

Intra-channel attribution looks at touchpoints on the same channel.

How Does Online Attribution Work?

There are several ways to identify which online channels and interactions a user engages with during their customer journey.

The simplest way is to use the referrer field in the HTTP protocol when a user is directed from an online channel to your website.

The referrer field is included with every browser request sent to a web server. Here's an example of a standard HTTP GET request:

```
GET / HTTP/1.1
Host: avenga.com
DNT: 1
Accept-Language: en-us
Accept-Encoding: gzip, deflate
Referrer: http://publisher1.com/article-about-adtech.html
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_12_6) AppleWebKit/603.3.8 (KHTML, like Gecko) Version/10.1.2 Safari/603.3.8
```

In this example, a user was reading an article on publisher1.com and then clicked on a link (or ad) and was directed to https://avenga.com.

Web Wanalytics tools, and AdTech and MarTech platforms display the following referrers:

Direct

When a visit is marked as direct, it means that the referrer information isn't known.

There are several reasons why a referrer can be marked as direct:

- The user manually entered the website URL into their browser's address bar or accessed it from their bookmarks
- The user entered the website from a subdomain, e.g., they first accessed publisher1.com and then clicked on a link that took them to blog.publisher1.com
- The user clicked on a link or ad in a native mobile app that didn't contain UTM parameters in the URL (e.g. publisher1.com/?utm_campaign=native-app)
- There were some technical issues that resulted in referrer loss, such as clicking on a link from a secure website (https://) to an unsecure website (http://)

The table below illustrates when referrer loss occurs between secure and unsecure websites:

HTTP Protocol	Referrer Passed or Lost?
https:// to http://	Referrer lost
http:// to https://	Referrer passed
http:// to http://	Referrer passed
https:// to https://	Referrer passed

Since most websites now use the secure https:// protocol, referrer loss between sites is much less of a concern.

Organic

Organic traffic originates from search engines such as Google Search, Bing, and DuckDuckGo. If a marketer is running *paid search ads*, those visits are typically classified as campaign traffic rather than organic.

Social

Visits from social media sites like Facebook, LinkedIn, X, and YouTube are marked as 'social'.

Website

When a user clicks on a link from a website and is directed to the advertiser's website, then it's classed as a 'website' referrer.

Campaign

The 'campaign' referrer is recorded when the website a user lands on contains UTM parameters. In this case, the Referrer field is ignored and the UTM parameters are used to determine the referrer. The 'campaign' referrer is recorded when the website a user lands on contains UTM parameters.

In this case, the referrer field is ignored, and the UTM parameters are used to determine the referrer.

Some AdTech and MarTech platforms simply classify these visits as campaign, while others provide more granularity — labeling them as paid social or paid search, depending on the information in the UTM parameters.

For example:

- If a user clicks on a LinkedIn ad and lands on a page with the UTM parameters shown below, the referrer might be recorded as *campaign* or *paid social*.
- If a user clicks on a Google paid search ad and lands on a page with the UTM parameters shown below, the referrer might be recorded as *campaign* or *paid search*.

Each time a user arrives from a different channel, a new session begins, and the referrer's information is captured — helping advertisers and marketers reconstruct the full user journey across touchpoints.

Online Attribution Models

1. The Last Click Attribution Model

The last-click attribution model (also known as last interaction or last touchpoint) is the oldest and still the default model in many web analytics, MarTech, and AdTech platforms.

In this model, 100% of the credit for a conversion goes to the final known referral, click, or traffic source before the conversion occurred.
For example, if a user's final action before converting was a direct visit, then the conversion would be fully attributed to that direct entry.

Although this model is one of the simplest, it ignores all the other touchpoints in a customer journey, which can lead to poor decision making when choosing which channels to optimize.

2. The Last Non-Direct Attribution Model

The last non-direct attribution model is very similar to the last click attribution model, but it removes direct visits from the equation.

In this model, 100% of a conversion is attributed to the last known referral that wasn't a direct visit.

Here's an example of how this process looks:

1. A user clicks on a link on Facebook and is directed to your website
2. They browse your website but then leave
3. They later type your website into their address bar and download one of your eBooks

Because the third step is a direct visit, the last non-direct attribution model would ignore this and assign 100% of the conversion to Facebook.

This model provides a slight improvement over the last-click approach but still overlooks the other touchpoints a user interacts with throughout their journey. As a result, it can lead to incomplete insights and poor optimization decisions.

3. The First Click Attribution Model

The first click (aka first interaction or first touch) attribution model is similar to the previous two, except it assigns 100% of a conversion to the first click or referrer in the customer journey.

4. The Linear Attribution Model

The linear attribution model evenly attributes conversions to all touchpoints in a customer journey.

Even though this model values each conversion equally, which rarely is the case, it is useful for getting an overview of the customer journey.

5. The Time Decay Attribution Model

The time-decay attribution model is a variation of the linear model.

In this model, the touchpoint closest in time to the conversion receives the most credit, while earlier interactions are assigned progressively less credit. Put simply, the farther a touchpoint is from the conversion event, the more its influence "decays."

This approach not only provides a fuller picture of the customer journey but also applies weighted importance to each touchpoint based on its proximity to the conversion.

The above image illustrates how different touchpoints could be attributed to the conversion.

This model assumes that the most recent touchpoints were the ones that influenced the user to convert, which may or may not be the case.

6. The Position Based Attribution Model

The position-based attribution model grabs all the touchpoints in the customer journey and assigns them based on their position in the conversion path.

This model gives more weight to the first and last interaction in the customer journey. The rest of the attribution credit is divided among the remaining touchpoints.

This model is often a good choice for advertisers as it provides an overview of the customer journey and assigns credit to the two most important interactions — the first and last interactions.

7. The Custom Attribution Model

Some AdTech and MarTech platforms allow you to create custom attribution models whereby advertisers set their own rules for attributing touchpoints in a customer journey.

This flexibility allows advertisers to account for their unique campaign structures, audience, and customer journeys.

However, the attribution models discussed above apply only to a single device and web browser (e.g., a laptop on Chrome or a smartphone on Safari). To accurately measure interactions that occur across multiple devices or browsers, advertisers need to implement cross-device attribution.

Cross-Device Attribution

Cross-device attribution aims to record interactions a user has with a brand across multiple touchpoints and devices, and ultimately attribute conversions accordingly. Where online attribution models aim to attribute conversions across different channels, multi-touch attribution aims to attribute conversions across different web browsers and devices, as well as channels.

Here's an example of how cross-device attribution would look:

How Does Cross-device Attribution Work?

To attribute online conversions across multiple channels, AdTech and MarTech platforms traditionally rely on cookies — typically third-party cookies.

However, since cookies are tied to a single device and browser, they can't transfer between environments, making them ineffective for cross-device attribution.

To bridge this gap, measurement platforms use deterministic, probabilistic, or hybrid matching methods:

- **Deterministic matching** relies on common identifiers such as email addresses or phone numbers to accurately recognize and link users across devices
- **Probabilistic matching** uses non-unique signals — like IP addresses, device characteristics, and location data — and applies algorithms and statistical modeling to estimate user identity

AdTech and MarTech companies create user profiles containing these identifiers and signals, then connect them through an identity or device graph — a data structure that maps relationships between users, devices, and browsers.

For walled gardens, cross-device attribution is significantly easier due to their deterministic data (login information, account IDs) and the fact that users often stay logged in across devices.

By contrast, independent AdTech platforms, brands, and agencies rely on DMPs to unify fragmented data. These platforms aggregate user information from various online and offline sources, build profiles, and construct identity graphs used for targeting, measurement, and attribution.

Offline-Online Attribution

Even as advertising and marketing continue shifting toward digital channels, there remains a strong need to connect offline and online data — enabling advertisers to track and attribute user actions across both worlds.

For instance, understanding whether a billboard campaign or radio spot led to a website visit or online purchase provides valuable insight into cross-channel performance and ROI.

Examples of offline channels include:

- Direct mail
- Traditional outdoor advertising — i.e. out-of-home (OOH) and digital out-of-home (DOOH) advertising.
- Telemarketing
- TV
- Radio

Below are the most common ways for advertisers to attribute offline ad exposure with online conversions.

Vanity URLs

Vanity URLs are domain names that are often created for a specific advertising campaign. They're designed to align with a company's brand, be easy to recall, and often redirect to a longer or more complex landing page URL.

Marketers commonly use vanity URLs in OOH, TV, and radio campaigns to make it easier for audiences to remember and visit the promoted site — especially when users can't simply click a link.

For example, instead of using a URL like *company1.com/new-product?utm_source=ooh&utm_medium=billboard-airport&utm_campaign=new-product*, an advertiser could use newproduct.com.

The vanity URL could take the user to a dedicated landing page, e.g., newproduct.com or redirect them to a different landing page, e.g., company1.com/new-product. With either option, the vanity URL will redirect the user to a destination page and add campaign tracking parameters for traffic attribution.

There are different types of vanity URL:

- Standalone vanity URLs, e.g., newproduct.com
- Subpage vanity URLs, e.g., company1.com/new-product
- Shortened vanity URLs, e.g., sv.ly/newproduct

Vanity URLs help measure the reach and impact of offline ad campaigns by linking them to online visits and conversions. However, they are not completely accurate, as some users who see the ad may instead search for the product or service on Google rather than typing the vanity URL directly. In that case, conversions would be attributed to Google Search instead of the offline campaign.

Despite this limitation, vanity URLs remain a valuable way to gauge the effectiveness of offline advertising efforts.

Time-Limited Attribution Windows

Another way to measure offline ad exposure to online web traffic and conversions is via the time-limited attribution window.

This model analyzes the period of time, e.g., 30 minutes, after the airtime of a TV or radio advertisement and looks for increases in web traffic and conversions.

When applying this model, advertisers need to consider the following:

- Window duration — How long should the attribution window remain open? For instance, should you analyze traffic and conversions within 30 minutes of the radio ad airing, or extend it longer?
- Traffic isolation — How can you distinguish between visitors influenced by the campaign and those who weren't exposed?
- Campaign overlap — How can you determine whether other marketing activities contributed to the observed increase during the attribution window?

While most AdTech and MarTech platforms provide attribution-window functionality as part of their measurement capabilities, it's typically limited to individual channels, e.g., display. To measure the online impact of offline campaigns, advertisers must either configure this manually within their analytics tools or use a dedicated cross-channel attribution solution.

Online Surveys

Instead of using complex attribution models, advertisers could simply ask users how they found their website.

Although this is a very simple approach, it can provide valuable insights that you might not get with attribution models.

Advertisers can implement online surveys at three different levels:

- When a user fills out the purchase or sign-up form or on the confirmation page
- When a user is browsing your website, open a discreet sidebar pop-up asking them to fill out a survey (you could offer a coupon code as an incentive)
- When the user is leaving the website, open a pop-up survey

Coupons

Coupons have been around for decades, but their popularity and effectiveness for attributing conversions is still strong.

By using coupons in their marketing materials, advertisers can attribute conversions to specific offline channels, often with more accuracy than with attribution models and technology platforms.

Coupons work best with direct mail campaigns and other printed advertising materials, but it's a good idea to issue unique coupons per campaign, and when possible, per client.

Zip/Postal Codes

Collecting ZIP codes from online customers can be used to measure different offline campaigns, such as direct mail and out-of-home campaigns.

While this method alone isn't highly accurate — you can't be certain that someone from a particular ZIP code was influenced by an offline ad — it becomes more valuable when combined with other attribution models for cross-validation and improved precision.

This approach is most practical for ecommerce businesses or companies with both online and physical stores, since they naturally collect billing and shipping information during the purchase process.

Online-Offline Attribution

Let's now look at a few ways advertisers can attribute online activity, such as ad views and clicks, to offline purchases in a store.

Beacons

Beacons are Bluetooth-enabled devices that can transmit signals to and from mobile devices such as smartphones and tablets.

When placed in brick-and-mortar stores, they can be used to send push notifications to devices in a certain radius and collect data about the device itself. The latter can help attribute online activity, such as ad clicks and mobile app activity, to offline purchases.

Zip/Postal Codes at POS

You read earlier about how advertisers can use zip codes to attribute offline ads to online conversions — but the process can also work in reverse.

One of the most common ways to collect zip codes from customers is by asking for them at point of sale (POS).

Advertisers can then match the zip codes collected in-store with the location data from their online ad campaign reports. However, just like when using zip codes for offline-to-online attribution, this method is not highly accurate and is best used as a supplementary approach alongside other attribution methods.

11. The AdTech Dilemma: Build vs Rent vs Buy

When it comes to selecting an AdTech platform, there's no one-size-fits-all answer. Businesses must decide between renting, buying, or building — each with unique benefits and drawbacks.

Build

Building means developing a custom AdTech platform from scratch or replacing components in phases. At Avenga, we know this topic inside out. If you want to create from scratch, refactor or partially improve your solution, we can help you bring it to life.

Advantages

- Complete control — Platform designed to your exact workflow
- Future-proof — Scalable alongside your growth
- Cost becomes investment — Ownership turns monthly fees into long-term asset value

Challenges

- Time-consuming — Developing a full MVP can take 3–6 months
- Higher upfront investment — Requires a skilled, dedicated engineering team
- Scope management — Risk of overcomplication

Rent

Renting means leveraging software-as-a-service (SaaS) platforms — perfect for companies at the start of their AdTech journey.

Advantages

- Rapid time-to-market — Minimal setup, easy onboarding, and low upfront costs with pay-as-you-go flexibility
- Learning period — Gain clarity on your real needs before committing

Challenges

- Rising costs — Expenses grow as usage and data volume increase
- Limited customization — Features and workflows can't be fully tailored
- No control over roadmap — Development depends on vendor priorities
- Transparency concerns — Restricted visibility into platform operations and data handling

Buy

Buying involves acquiring an existing company and its technology — including the product, codebase, intellectual property, and often the development team behind it.

Advantages

- Full ownership and transparency — Control over features and data
- Accelerated timeline — Faster than building a platform entirely from scratch
- Control over roadmap — Adapt and evolve the acquired technology to your business strategy

Challenges

- Integration complexity — Merging the acquired system can be time-consuming and resource-intensive
- Retention risks — Key team members may leave post-acquisition, impacting knowledge transfer
- Inherited issues —Such as technical debt, architectural limitations, and unresolved bugs

	Rent	Buy	Build
Time-to-Market	Fast (days/weeks)	Moderate (weeks/months)	Slow (3–6 months for MVP)
Upfront Cost	Low	Medium to high	High
Ongoing Costs	High (scales with usage)	Medium	Low to medium (maintenance)
Control Over Features	Low	High	Very high
Scalability	Limited by vendor's roadmap	Flexible after integration	Fully customizable
Transparency	Low	High	Very high
Integration Complexity	Low	High	Medium
Risk	Low initially, higher long-term	Medium	High initially, lower long-term

Key Takeaways:

1. Start with rent
— It's ideal for fast launches and early learning

2. Upgrade to buy or build when ready
— If your tool can't adapt to you, it's time to change

3. Focus on core needs
— You likely only need 2–5% of big platforms' features

4. Incremental builds work
— Replace third-party parts gradually

5. Custom doesn't mean complex
— Narrow scope ensures faster delivery and ROI

6. The tool should serve you
— Not the other way around

Your choice between rent, buy, and build will depend on:

- Stage of business maturity
- Budget
- Urgency to launch
- Long-term strategic control needs

If you start small, learn fast, and upgrade strategically, you'll avoid costly missteps and end up with an AdTech stack that truly fits your business.

Final Words

Thank you for learning with our AdTech Book!

If you have any questions about the contents of this book, or simply want to provide feedback, then please contact us via adtech@avenga.com.

Best regards,

Avenga AdTech Team

avenga.com