

Chapter 3 - Section 8

Affiliates and Affiliate Networks

- Affiliate Sites and Networks
- How the Affiliate Earns
- Auditing the quality of the Long Tail
 - *Figure 3.8.1*
- Tracking Affiliate Activity in third party ad serving
 - *Figure 3.8.2*
- The absence of impression tracking
- Affiliates that grow

This section will look at Affiliate advertising as a channel and put it into the context of third party ad serving. Like Search, ad servers have a greater role to play in tracking rather than serving the physical ads and so the benefits of a marketer, Advertiser and Agency using a third party ad server in this space are the same as apply to the world of paid and natural search. We will begin by describing what Affiliate advertising is before diving into the complexity behind the tracking methodologies required and the reporting capabilities that this provides. Being a performance marketing channel, Advertisers which are considered to be Direct response Advertisers are more common users of the Affiliate Channel than those seeking to do big branding exercises. This is because Affiliate activity at scale can be a cost effective method to drive sales and other conversion activities (see Chapter 4).

Affiliate Sites and Networks

The landscape of Affiliate advertising as a digital marketing channel is made up of two core entities: the Affiliates websites and the Affiliate networks which house them. Typically speaking the Affiliate network supplies the digital creative (core images and *.swf files) and click trackers to a mass of Affiliate websites.

Large Advertisers, particularly in the gambling and online retail verticals, have their own Affiliate programs so do not always utilize Affiliate networks. Instead act as the network themselves by storing the creative and distributing the click trackers. In doing so, the Advertiser takes on the responsibility to provide the Affiliate websites with their reporting data.

Affiliate sites go on to earn a commission when users click the affiliate links and purchase on the Advertiser site right away or within a given time frame. It is in the best interest of the network to drive traffic that will end with a purchase, so often the Network will encourage affiliates to use the context of their digital publications to incentivize the user to click on the ad. Typically speaking eighty per cent of all traffic is generated by Affiliates that offer an incentive scheme. That is that they offer a discount, a general offer with money off, or some other type of monetary incentive such as virtual points. The remaining twenty per cent of Affiliates tend to keep themselves content focused such that the already niche nature of their websites utilizes the SEO framework of the site design, and this in itself attracts users. The user searches a relevant term in the search engine, finds the Affiliate site perhaps halfway down the page of the search results and clicks to the Affiliate site. Here they find content of interest and a contextually relevant ad, which drives a click and potentially a sale.

How the Affiliate Earns

Networks such as *Trade Doubler* and *Commission Junction* have gradually attracted the attention of smaller Publisher sites that sign up to the network to become Affiliates. Within the interface supplied by the network, the website manager (for the now Affiliate website) browses the ads available to the Affiliates on the network and enrolls in the campaign. Enrollment releases the ability for the Affiliate to implant the ad on the site. Clicks on the ad are usually counted by the Affiliate network and are displayed

in that interface alongside attributed conversions (see Chapter 4) so that Affiliates can determine their daily revenues. This method does not work for all small Publishers, since the Publisher is essentially displaying the ads for free and only getting paid when the user clicks or performs a conversion activity after clicking. This can result in the Publisher giving away millions of impressions for free with no return for the use of the ad space.

Auditing the quality of the Long Tail

It can be very hard for a Network and Advertiser to audit these small Publisher sites on a continual basis, which can result in some negative consequences for brand image. Generally speaking an Affiliate site will be tailored towards the purchase activity and therefore will tend to abide by an Advertiser's branding regulations which should be supplied by the Network in the enrollment phase. This lack of control has put some Advertisers off of using affiliates but new developments in verification reporting in the third party ad server are likely to remedy this over the next few years (see Chapter 5 - Ad Verification).

Without using technology solutions there are checks that can be done on the Affiliate sites simply using the operational headcount available either at the Media Agency, the Affiliate network or on the Advertiser's marketing team. There are two processes for policing: "Affiliate Auditing" & "Affiliate Screening"- both of which would be outlined to the Affiliate in the enrollment phase.

Auditing is usually a quarterly process that involves looking over every individual ad that is being shown across all sites and checking copy and ad position to see that it aligns to the Advertiser's branding regulations.

Screening is a process of accepting or declining affiliates who apply to join the program. The advertiser or affiliate network can decide which affiliate is good enough to join the campaign and thus can stop low quality sites from enrolling.



Figure 3.8.1 *Affiliation extends the distance between the planners and buyers, and the inventory that is being procured. To ensure that the final location of the ad is safe for the brand to sit inside there are several safety nets both operationally and technologically. The auditing and screening process is covered in detail here and Verification is covered in depth in Chapter 5. It is worth noting that verification works off of the impression rather than the click metric on which most affiliation is based, so thorough verification reporting is not available for campaigns conducted across a click-based buying model.*

Tracking Affiliate Activity in third party ad serving

In the third party ad server the trafficker will see the name of a whole Affiliate network on the media plan in place of a Publisher site name. In the third party ad server interface, in the usual trafficking workflow for Display activity, the affiliate network name can be selected under the ad server site names. Once the site is added to the campaign in the interface, the trafficker generates a single tracking placement for each Affiliate network listed. Only the click portion of the placement is required and when the tag is trafficked to the affiliate network, it is this portion that the Affiliate network operations team load onto the enrollment page of their network; representing the Advertiser's campaign. The Affiliate network then adds information to the page about brand regulations and in its completion the package is

known as an "affiliate program". When the affiliate logs into the interface of the affiliate network, they enroll in the program. Once the ad is on their site, users that click the activity will trigger a count for a click in the third party ad server.




Figure 3.8.2 Affiliate tracking consists of a single tag or click tracker provided by the trafficker in the ad server interface: supplied to the Agency or affiliate network operations team directly who will set up the tag to load the affiliate identifier into the click tracker as a token when the tag fires. When the user arrives on the Affiliate site, the identifier for that site is passed into the token in the single network tag as the click data is passed into ad server reporting. The user then continues on to the landing page and the conversion data (see Chapter 4) is also captured in the ad server reporting. Within reporting a feed is compiled that shows what activity (clicks and conversions) should be attributed to each affiliate which can include the monetary value of the conversion or total conversions. This information is then used by the Affiliate Network to pay the affiliates.

Since the Affiliate network mostly pays out on a CPA and not a CPC, the trafficker in the third party ad server interface also needs

to ensure that a conversion tag is generated and is implemented on the Advertiser site in the right place (more on this in Chapter 4).

The absence of impression tracking

Advertisers do not use the impression portion of the ad tag in their Affiliate activity because the affiliate program is usually opened up to an unknown number of Affiliates, who will deliver an unknown quantity of impressions of the ad in order to get clicks or conversion events. Affiliates are commonly asked to house the creative or images on their own servers so that even the Affiliate network is not burdened with serving the actual creative. If the impression portion was used or the ad served, the ad serving costs to the Advertiser could be astronomical since it is a very broad base of long tail Publishers driving the traffic to the Advertiser site. It would also be very difficult to contact all the Affiliate sites to ask them to take down their impression pixels to cease ad serving entirely. For this reason is not effective for brand advertiser to use the Affiliate channel for brand activity since clicks and conversions alone are not a good representation of true reach and frequency (see Chapter 5 - Reporting).

Affiliates that grow

Affiliate activity does not just consist of small Publisher sites. Affiliation is a very clever and complex dynamic of shrewd Publishers that use very innovative site designs in order to drive user traffic to Advertiser sites for the purpose of purchase. All such Affiliates start out small but their content is arranged in a method that attracts high volumes of traffic. These affiliates are sometimes referred to as "Aggregators" and often break away from Affiliate Networks to become large Publishers in their own right, attracting Advertisers and Agencies to spend with them directly.

Just as with Search activity, the Advertiser or Agency will just end up paying the ad serving provider based on the click activity. To

conclude this chapter we will turn now to the outlying channels of Digital Advertising. From an ad serving perspective there are a number of restrictions placed on these channels which severely limit the existing tracking and serving capabilities of ad serving technology. Such restrictions exist not because the Advertisers do not want to track their activity but that Publisher-side technologies place restrictions on the ability for a true independent audit to occur. These channels will continue to suffer from funding problems unless the Publishers open the doors, and offer greater transparency and a hand of integrity to those with the marketing budget to spend across a wide reach of channels.

Chapter 3 - Section 9

Email, in-game and Social

- Permission Marketing
 - *Figure 3.9.1*
- Sending Emails on Mass
- Email Marketing advertising metrics
- Third party ad server tracking for Email
- In-game and digital outdoor tracking
- Social Media tracking
 - *Figure 3.9.2*
- The versatility of tracking pixels

This section rounds up the email, in-game and social channels for digital advertising. These channels are not small channels by any stretch of the imagination, they get a lot of user attention and vast volumes of internet traffic through them every waking moment of the day. However when trying to use third party ad serving as eyes across all media buys, these channels are outliers and for one reason or another limit the amount of tracking and serving that can be done. Thankfully the meagre technology of tracking pixels and click trackers allows a very independent eye on delivery and click activity, but beyond this Advertisers are essentially blind and must submit to Publisher reporting at the very best.

Email advertising is an old channel in the world of digital advertising. Few Advertisers extend their budget to cover it beyond their own existing user base but if it is a traffic driver to an Advertiser site, users need to at least be tracked for the email channel to appear in the ad server reporting. In-game is classed as standalone software either on a PC or secondary device including console. Along with Email, In-game as a channel suffers the same

setbacks as trying to track activities with applications for Mobile advertising. Lastly we will look at social: a dominator in the world of content production but very slow to monetize their properties. Social networks continue to keep their networks locked down to third party ad serving which seems daft, given that it shuts off potential revenue in favour of less transparency for the Advertiser only adding to the speculation that social may not be as effective for Advertisers as initially thought.

Permission Marketing

Email advertising gets a bad rap because Advertisers believe that email advertising is synonymous to Spam. Strictly speaking the difference between email marketing and spam is that the user has provided consent to be sent emails from a third party versus being delivered unsolicited mail. The difficulty with presenting email marketing as a fruitful channel for Advertisers is that the methods to obtain consent for email marketing are lumped into small illegible boxes as a condition to access a service. With such methods the consent is provided but the user never really wanted to receive advertising from a third party.

Email advertising falls into two broad categories: emails that are single page adverts, and monthly subscribed emails from a Publisher directly with ad space around the sides or embedded among the content. To receive either of these two forms of email advertising, a user must sign up to the service and opt in to receive messages from third parties. The strict opt in requirements demanded by regulators such as the DMA (Direct Marketing Association) in the UK mean that the email advertising channel is sometimes known as "Permission Marketing" a phrase also coined by the author *Seth Godin*.

A diagram of a web form. At the top is a wide rectangular box. Below it are several horizontal lines representing text input fields. Two of these lines are labeled with numbers '1' and '2' in small black boxes on the left margin. At the bottom right of the form is a button labeled 'Submit'.

Figure 3.9.1 *The cursed double-opt-in boxes. On a Publisher site a user might be interested into the content as to opt in to a newsletter. On the newsletter opt in page there may be displayed two boxes: (1) is for first party advertising: can the Publisher deliver Publisher-specific ads to the user's inbox? The second (2) asks about third party advertising, upon which email advertising is founded. Publishers regularly try to use the manipulation of language to confuse the user into opting in. Examples include "tick here to not opt out of third party advertising"; which will only serve to de-value email as a marketing channel even further.*

By consenting to receive such emails, the user's email address becomes a commodity that the Publisher or email list owner can sell to Advertisers or, with a user's permission on a second tick box, the "double opt-in" can be sold on to other list owners to receive additional mail where further advertising space is then made available to Advertisers.

Sending Email on Mass

Consented emails from respected Publishers can be commoditised because the emails are opened and read and the ads attract clicks, which can result in conversions (sometimes with

better results than standard display - See Chapter 4). In order for a Publisher to send thousands of emails they either use their own list or buy one in from a data provider such as *Axiom*. As such the list of emails becomes the property of the Publisher for a given period of time and it is their responsibility to act in accordance to the wishes of the Advertiser, the user and the original list owner.

Email list quality is a huge factor determining the success of a campaign and it essential that addresses are cleaned regularly by the database owner to ensure that any addresses requesting an unsubscribe or where a "bounce back" is occurring, are removed. To send thousands of emails, a Publisher needs to use a mass email delivery system such as *Email Vision*. The list is simply uploaded to the delivery platform, the HTML email compiled, and then sent as an email over a given period of delivery. To be certain the emails are not blocked, the Publisher needs to be sure they are sending from a whitelisted IP address. They can work with the email delivery platform and major Publishers like MSN or Yahoo! to check that the IP is not on a blacklist where other spam email sender addresses are blocked. Here, the email provider is acting as the gate-keeper for what can and cannot be loaded into emails. If an IP address is deemed dangerous, links or images can be blocked and in some instances emails can end up in junk mail folders.

Email Marketing advertising metrics

An email delivery platform can measure four core metrics.

1. Open rate, detecting how many of the emails sent were viewed in the main reading pane of the email program or web-based application.
2. Click through rate (CTR) for the subsequent clickable areas or URLs embedded in the email.

3. Bounce rate (i.e. the number of emails that were never delivered to inboxes because they received a bounce-back). These email addresses are sent back to the list owner after the campaign ends to investigate to see if they are now wholly inactive.
4. Unsubscribe rate; needless to say a high unsubscribe rate following the launch of an email campaign, is a storm cloud that neither Advertiser nor Publisher wishes for. The user that clicks the unsubscribe button, which should appear at the bottom of every Publisher-owned email, will be removed from the original list (which will either be Publisher owned or hired).

Third Party ad server tracking for Email

Connecting the third party ad server to the email marketing channel is important to ensure that this channel is independently tracked and that the email channel is considered as an influence in the user's pathway to conversion (see Chapter 4).

The trafficker in the third party ad serving interface would look on the media plan and might see a Publishers name followed by a mention of the name of a newsletter. This same reference would need to be included in the third party ad serving site list. The traffickers selects this as a site and generates a 1x1pixels and click trackers and sends these on to the Publisher ad operations team to ensure they embed them into the email creative. True third party ad serving (even of images inside emails) is never undertaken. More often than not, the impression tracker is also wholly abandoned and only the click tracker is implemented. Here, the email provider is considering the third party ad servers IP address to belong to a blacklist because there is such a wide usage of the IP across millions of other emails at some point in the past. This means that pixel never loads and therefore never records an impression.

In addition not all email provider interfaces are setup to receive emails that can render images properly. Those that are not 'HTML friendly' will render a text-only version of the email. In such a delivery the image pixel would never load but the user may choose to click through via the supplied link, which would route through the third party ad server before landing the user on the landing page.

Due to the simple nature of email and email delivery there is no possibility of recording dwell rates or interaction rates, but conversion tracking is possible so long as the click trackers are in place on the email creative (more in Chapter 4).

Email lists have evolved significantly in recent years. Formerly they were large databases of Publisher owned data about users where users were recognised by an email address collected with their permission. Now list owners and Publishers have gone a step further and logged the cookieID of the users using first party cookie information and have also labelled up the users into audience types by logging their interests based on what they are being exposed to and what they are clicking on. Rather than sell these databases to Advertiser's as email lists, they are now being sold as sources of data to be inserted into DSP's and DMP's to improve programmatic buying and trading (see Chapter 8).

In-game and digital outdoor tracking

Despite the promise of mass in-game advertising, high reach games remain a premium offering to players and therefore advertising rarely has much impact in this space. In game ads would deliver creative to an online game via a feed (such as an XML feed). The same is true of digital outdoor advertising as well (content can be kept regularly up to date via an XML feed) but the nature of the games console and standalone gaming applications outside of the browser, mean that the third party ad serving cookies cannot be dropped. This hinders the ability for Advertisers

to establish how much value the game environment really provides to the overall cross-channel ad campaign. Until such time as in-game cookies are dropped and can easily be matched with those of the ad server; there is less strength in running such campaigns as the impact of the creative itself in the context of a multi-channel campaign is lost. There may be some hope in second screen audio watermarking technologies but at time of writing, nothing scalable has come onto the market.

Social Media tracking

Digital advertising through social channels is still in its infancy and as yet the link between third party ad serving and these environments is rather tentative. Social media in the context of Digital Advertising can be divided up into three parts:

1. The mass advertising that appears on the right side of the page for Facebook. This is managed by an advertising exchange program called FBX (see Chapter 8). This scaled technology is much more closely aligned to the world of Paid Search than that of display advertising.
2. Ads that appear in the user feed. These ad spots are technically unlimited and depend on the user engagement with the private feed. Ads are purposed from content on fan pages or elsewhere on the web but have limited engagement capabilities.
3. The home of the Advertiser on the social network. In the case of *Facebook*, this is the fan page, where it has been notoriously challenging to embed tracking and tagging technologies such as conversion tags (see Chapter 4). Meanwhile 'tweets' and 'likes'; the measures of user engagement endorsed by the social networks Twitter and Facebook, do not pass this information back to the third party ad server using any reserved method (so officially the metrics are not supported by third party ad servers).




Figure 3.9.2 *Still in its infancy: the ad ecosystem of the popular social networking site Facebook consists of 3 core properties for Ads. FBX (which caters for the paid scaled and targeted advertising on the right navigation: drawing from user data), the Advertiser's Fan Page (for some Advertisers this is more of a hub than their own websites) and ads which appear in the user's news feed (commonly content is repurposed from the Fan Page).*

In all of these situations one fact rings true over all others as a cry of utter restriction. Facebook and Twitter do not allow JavaScript tags from an external source on their sites as a matter of policy. This means that should these sites appear on the media plan, the trafficker merely sets up impression and click trackers for each placement. The trafficker then sends the resulting tags along to the Media Agency to either embed into the FBX access point or work with technology providers that are repurposing Advertiser content as ads (such as *Upcast* or *Experian's Alchemy Social* product suite). This data leaves very limited performance information (clicks and conversions) and delivery information (impressions) inside ad server reporting for the Advertiser and Agency to see.

The versatility of tracking pixels

It is clear that tracking pixels and click trackers are very versatile technologies and allow for tracking in a variety of channels including a majority of the newest channels available to Advertisers. The concern with pawning out a vast quality of trackers and not conducting full ad serving is that there is only the confirmation that a pixel was ever delivered on a site and that a click was instigated. There is no proof as such that the true ad was delivered or who saw it. Publishers that today only accept tracking pixels will need to go a step further and provide greater transparency to third party ad servers and technology providers if they are to attract full investment from Advertisers with big budgets.

This concludes the Channel by Channel campaign setup chapter of this book. We have gone meticulously through the work of the trafficker in the setup of ad campaigns via the third party ad serving interface, from Display, Rich Media, DCO, Instream Video and Mobile, to Paid Search, Natural Search, Affiliates, Email and the Social channel. But the campaign setup for all channels by the trafficker is only complete when the trafficker generates not just the ad tags but also the conversion tags. The next Chapter looks in depth at the importance of Conversion Tags and how the trafficker uses the concept of conversion attribution to discern the optimum targeting and audience settings available in the campaign setup.

Chapter 4

Attribution and Audiences

Chapter 4 - Section 1

Conversion tags and Attribution

- What are conversions?
 - *Figure 4.1*
- Where to place the conversion tag?
 - *Figure 4.1.1*
- Where to have conversion tags (table)
- CPA
- Conversion tag types - Sales Vs. Counter Tags
 - *Figure 4.1.2*
- Extended data/Extended Parameters
- *Figure 4.1.3*
- *Figure 4.1.4*
- How do you test conversions?
- Conversion Tag Functionality - Piggybacking
 - Container Tags
 - Certified container tag providers
 - Which conversion tag fires first?
 - Selective Pinging
 - *Figure 4.1.5*
 - Piggybacking with Extended data
 - *Figure 4.1.6*
- Conversion Tag Functionality - Conversion Windows
 - Window Settings
 - Post-impression and Post-click windows
 - Long and short conversion windows
 - Shadow Tagging
 - Reporting on Conversions
 - Repeat Conversion Windows
- Conversion Tag Functionality - Attribution Modelling
- *Figure 4.1.7*
- Attribution Settings
- Engagement Mapping

- *Figure 4.1.8*
- Custom Attribution Modelling
- Challenges in implementing conversion tags
 - Tag Management Solutions
 - *Figure 4.1.9*
 - Universal Tagging Security

In Chapter 3 we explored the campaign setup channel but all the efforts of the trafficker come down to one question for the marketer: What is the goal of the advertising campaign? What does success look like? For big brand clients the goal is usually to get message out to a target audience, to build awareness and brand affinity but the ultimate goal (as is the goal with more direct response style Advertisers) is to drive sales. This section will go into technical detail about "conversion tags"; a technology that can be used not just for measuring sales but also events that might be deemed a goal for the marketing team. Beyond this it is important to establish exactly what activity drove the use to a conversion; untangling all of their marketing engagements with an Advertiser to establish how to attribute the conversion to a paid-for activity.

What are conversions?

It's all very well measuring the delivery of ads and collecting information about how many unique users saw and clicked ads but the vast majority of Advertisers want to know how many users then landed on their website and how many went on to buy a good or service. This can be ascertained by monitoring the traffic driven directly to the purchase pages of the Advertiser website.

Conversions by definition vary by vertical. Here are some examples of uses for the conversion tag:

- Automotive - Tracking Brochure requests or test drive sign-ups
- CPG/FMCG - Tracking sample requests or coupon downloads
- Retail - Tracking sales through the site or clicks to the sales phone line
- Entertainment - Sales through to ticket purchase or local cinema listings
- Finance - Leads generated through Insurance quotes
- Telco - Form completion for a free SIM card

A conversion tag is a piece of JavaScript code; generated inside the third party ad server interface by the trafficker which is sent to the site management team of the Advertiser website to upload into the source code of the page. The trafficker can output the conversion tag in various formats (such as iframe Vs. JavaScript) depending on what is easiest for the Advertiser's website management team to implement.

Once implemented on the page, the conversion tag will allow the Advertiser to watch for incoming traffic and recognise it as originating from a third party served or tracked ad. The tag, once present on the webpage can recognise the difference between ad-driven traffic to itself which is called is a "Conversion", and non-ad driven traffic (which carries various names depending on which technology is being used; the *Atlas* third party ad server calls this an "Action").




Figure 4.1 Conversions detect for the presence of an existing third party ad serving cookie on a user's browser. This will have been dropped onto the user's browser by the delivery of an ad that the user would have encountered previously. The conversion tag then passes the cookieID to ad server reporting when it loads on the Advertiser site to allow ad server reporting to place the conversion event, marked with a timestamp to the cookieID in cookie level data. In aggregated reporting the Publisher which featured the ad would show a single conversion. In comparison if a user visits the Advertiser site directly, having not previously encountered an ad or having deleted their cookies, will still trigger the conversion tag on the page. This conversion tag will update the ad server reporting with an action instead of a conversion. This will add a record to the conversions and actions database in the ad server and not be marked or any cookie level data against any existing cookieID.

Once the "conversion tag" (also called a "Spotlight tag" in DoubleClick or an "Action tag" in Atlas) is generated and placed on a page it begins to work, feeding conversion and action information back to the third party ad server. It does this by looking to the users browser when they arrive on the page and once the tag code renders, seeks a cookie belonging to the third

party ad server; a cookie which would have been dropped by a third party ad server via an ad in the past.

If the conversion tag finds a cookie; the ad server matches the cookieID to the action of the conversion tag firing and in the ad server database adds a new row to a master table, to represent that the conversion occurred, when it occurred and which cookie it occurred on. If no cookie is found, the ad server still makes a log that the conversion tag was triggered and at what time, labelling it as an Action.

Where to place the conversion tag?

The Advertiser is advised to place a conversion tag on whatever pages relevant to measure the success of the campaign. Put a conversion tag on a landing page and the ad server can measure the number of users that got through to the Advertiser site. Place one deeper on the pathway to an actual purchase, such as the basket page, and more information is gleaned such as "drop-off rate" which indicates how many users are not making the journey all the way through to a purchase.

The most useful position for a conversion tag is on the "confirmation" or "thank you" page of an e-commerce site. When a user triggers the conversion tag on the confirmation page; the Advertiser knows that the user has purchased; a purchase which can be attributed back to the ad which drove the user to the site in the first place. It is imperative therefore that each conversion tag is unique to its page 'type'; placing a conversion tag on a landing page and then duplicating the code to use on a confirmation page will skew results. This is because the arrival on the two different page types could be counted as marketing success goals but they need to be segregated since an arrival is a less valuable "conversion event" than a purchase.




Figure 4.1.1 The paid-for media channel/Publisher site, drives the user through to the product landing page via an ad. There could be a conversion tag on this page (C1) to assess how many users get this far. Users that make it as far as the basket page could have this conversion event tracked by an additional conversion tag (C2) and those users that purchase will trigger the last conversion tag shown (C3). In ad server reporting it is possible to establish the drop-off of users throughout the process and correlate the difference between users that for example, came in off of a predominantly-social pathway versus those that came in from more of a display pathway. The results might indicate that users that see more display ads rather than social ads are more likely to cause C3-type conversions. Site side analytics can also measure such activity but the ad server factors in the user's engagement with all ads in the campaign to establish an attribution pathway (covered off later in this section).

The pages on the Advertiser site soon become markers in the marketing funnel for the ad campaign as users arrive to the site. Those that drop off at product pages are conducting research. Those that drop off at the basket page are closer to the buying decision. This progress should be reflected in the naming convention adopted by the trafficker when the conversion tags are created so that it is much clearer in reporting how valuable certain conversion events are.

Conversion tags are also useful in other locations:

<i>Where to have conversion tags</i>	<i>Why it is useful</i>	<i>Why it may not be possible</i>
Within ads	If the landing page is not accessible, the click within the ad can be used to trigger a conversion tag. If the conversion event is to fill in a form: the form can exist in the ad itself and the submission of the form can trigger a conversion tag.	Conversion tags can be used to carry retargeting tags and if these are fired in locations on Publisher pages then there is the opportunity for an Advertiser to build a cookie pool from the Publisher's audience (see Chapter 7 - Retargeting). This is a condemned practice called Publisher Data Leakage and should be avoided to maintain good relationships with the Publishers.
Facebook Fan Page or Youtube Channel page (On Advertiser	Brand Advertisers in particular drive traffic to <i>Facebook</i> fan pages or <i>Youtube</i> pages.	<i>Facebook</i> and <i>Youtube</i> (at time of publication) do not allow third party conversion tags on their properties. Advertisers will have to

'properties' outside of the Advertiser site)	A conversion tag here would show the value of the Publishers and ads driving the traffic.	work hard to demonstrate the importance of these technologies to these big Publishers in order to have conversion tags placed.
Amazon or E-Retailer or Supermarket sites (On alternative purchase locations)	FMCG, CPG, Electronics and Telco Advertisers in particular allow users to purchase their products via other online sites. Conversion tags on these sites will show how many sales are being generated and which ads are driving those sales.	Amazon and large supermarket sites (at time of publication) do not allow third party conversion tags on their properties. Advertisers will have to work hard to demonstrate the importance of these technologies to these big Publishers in order to have conversion tags placed in basket and thank you pages.

CPA

The success of a campaign can be measured by the number of completed conversions but in addition the actual sales can act as a trigger to paying the Publishers. A CPA is a cost model that agrees that the Advertiser pay outs a percentage of the value of the final sale to the Publisher that drove the user to the confirmation page. In return the Publisher delivers however many impressions, and thus provides however much inventory will be required to meet the income goal returned through the conversions.

As an example: an Advertiser might need to sell ten thousand suits and have an Advertising budget of £100,000 to do this, but approaching a Publisher site, finds that the CPM is £1. This would mean that every tenth user that sees the suit ad would need to buy the suit through the website. The Advertiser might know that

the likelihood of this happening is very small; a conversion is more likely to work for one user in a hundred than one in ten. The advertiser could negotiate with the Publisher to agree instead to pay on a CPA of £1. This way the budget is used to sell all ten thousand suits and the cost is lower for the Advertiser.

Clearly CPA does not work for every Publisher, but will work for Publishers displaying ads to niche audiences for high cost purchases. In such an example the CPA can return more revenue than a CPM.

If an Advertiser site contains more than one product for sale, the conversion tag can be used to work out which product was actually purchased and how much it cost:

Conversion tag types - Sales Vs. Counter Tags

Conversion tags come in a few varieties; most ad servers have the capability of generating conversion tags with the sole purpose of counting the traffic or the number of conversions (called a counter tag). This can be compared to the generation of what is referred to as a "Sales Tag"; which can pick up more information about the conversion in addition to counting, such as value of the user purchase or number of items in the basket.

Technically speaking these tag types are very similar. The differentiation helps the trafficker to remember the purpose of the tag. This gives an indication of the tag location when campaigns are set up without consistent naming conventions.




Figure 4.1.2 *The trafficker is best off creating counter type conversion tags for pages where the Advertiser just wants to monitor the user's journey to that page. Sales tags are better created for pages where the most valuable conversion event is occurring such as the point of sale.*

Extended data/Extended Parameters

The sales type of the conversion tag can also be setup to capture specific information about a sale. This information is made up of three groups of parameters referred to as "Extended Data" or "Extended Parameters". In the collection of each of these parameters, the Advertiser site management team pass the parameter into the corresponding field in the conversion tag using a mechanism like a simple JavaScript code, before it delivers the conversion information back to the ad server.




Figure 4.1.3 As the user makes the purchase on the confirmation page; extended data parameters can be supplied to the conversion tag via a script, and supplied into ad server reporting.

The three main groupings of parameter are:

1. "OrderID": The OrderID is a unique field generated by the Advertiser upon purchase to differentiate between the sales.
2. "Metric Parameters"; number value associated with the sale. The most commonly collected are two parameter fields called "Quantity" and "Revenue". Quantity usually reflects the number of items in the basket and Revenue reflects the total basket value.
3. "String Parameters"; are available to add context to the purchase. Some ad servers call this information "Category fields". Category fields follow strict guidelines; the Advertiser is permitted to pass information about the sale such as the names of items in the basket or description of the items purchased, but it is strictly prohibited for the Advertiser to collect Personally Identifiable Information.

Some ad servers limit the number of Metric and String Parameters that can be collected but there is rarely a need for more than one hundred fields of data. Most campaigns utilising sales conversion tags will collect between three and six data points.

```
<script type='text/JavaScript'>
// Conversion Name: Confirmation Page
var ebRev = '%JS_TotalCost%';
var ebOrderID = '%JS_ConfirmNo%';
var ebProductID = '%JS_RmTypCd%';
var ebProductInfo = '%JS_HName%';
var ebRand = Math.random()+'';
ebRand = ebRand * 1000000;
//
document.write('&lt;scr'+ 'ipt src="HTTPS://bs.serving-
sys.com/BurstingPipe/ActivityServer.bs?cn=as&amp;Activity
ID=32932240&amp;rnd=' + ebRand +
'&amp;Value='+ebRev+'&amp;OrderID='+ebOrderID+'&amp;
ProductID='+ebProductID+'&amp;ProductInfo='+ebProductI
nfo+' "&gt;&lt;/scr' + 'ipt&gt;');
//]]&gt;
&lt;/script&gt;

&lt;noscript&gt;
&lt;img width="1" height="1" style="border:0"
src="HTTPS://bs.serving-
sys.com/BurstingPipe/ActivityServer.bs?cn=as&amp;Activity
ID=3299410&amp;Value=[Revenue]&amp;OrderID=[OrderID
]&amp;ProductID=[ProductID]&amp;ProductInfo=[ProductIn
fo]&amp;ns=1"/&gt;
&lt;/noscript&gt;</pre>
</div>
<div data-bbox="139 738 936 901" data-label="Text">
<p><b>Figure 4.1.4</b> An example of a Mediamind sales conversion tag. Such a tag is best implemented as close to the top of the source code on the Advertiser page as possible (to ensure the conversion event is captured if the user leaves the page before it is fully loaded). This script would be generated by the trafficker who would send it through the Advertiser website management team. There are a series of extended</p>
</div>
<div data-bbox="479 923 525 941" data-label="Page-Footer">345</div>
<div data-bbox="466 961 527 982" data-label="Page-Footer">346</div>
```

parameters that the trafficker has added to the example tag shown which include both Metric Parameters and String Parameters. Product ID is an example of the former and ProductInfo is an example of the latter. The tag shown has two portions: the part labelled <noscript> can be implemented on a page that does not accept JavaScript but only the <script> or the <noscript> portion of the HTML shown needs to be implemented at all.

It becomes part of the trafficking process for the trafficker to add the Metric and String Parameters as required to the tag before sending it over the Advertiser site management team with implementation instructions.

With regards to the collected data, the ad server simply records the data against a CookieID (which cannot be traced back to an individual's personal details). That CookieID is associated with an OrderID when the conversion tag fires and the detail are recorded in a single row in the ad server reporting database (see chapter 5). Here inside the confines of the ad server, no information at all about the purchaser can be gleaned except the behaviour of the cookie in question by looking in the database to other references to the Unique CookieIDs engagements with other ads or conversion tags. A consistent collection of information for a campaign from a conversion tag, will allow extended data parameters to become columns in ad server aggregated reporting (see chapter 5) and purchases can be grouped so that users that purchase a particular product, can be analysed for their buying patterns, ad engagement habits, expenditure and complimentary purchases.

Once the trafficker has exported the conversion tag from the ad server interface, adding metric and string parameters; the site management team will implement the tag. At this stage, testing out the implementation is crucial to avoid mishaps in data collection down the road.

How to test conversion tags

Conversion tags work right out of the box. Once implemented, they will count conversions in ad server reporting if a user's cookie has encountered an ad for the same Advertiser from the same ad server before. This is not universal; some ad servers allow conversion tags to only count conversions if users have encountered ads from a campaign associated with that conversion tag.

As an ad server begins tracking the conversion data, the reported figures should remain relatively consistent, however if ad server reporting shows a fall in conversions or if the conversions do not start appearing right after a campaign is put live, then the conversion tag itself needs to be checked.

An implemented conversion tag can be visited just by landing on the right page; if the user has not encountered an ad before, the tag will fire; triggering an "action" which will show up in ad server reporting. The user arriving to the page where the conversion tag sits can use an HTTP tracer (which, as a reminder, measures the server calls to and from the page).

If the URL of the ad server appears in the list of server calls when the page is loaded, then the conversion tag is present. To check its functionality deeper to determine if the conversion tag fired off an "action" versus a "conversion" requires to user to know what a successful conversion count from the conversion tag to the server looks like. This information can either be obtained by contacting the third party ad serving provider or in some instances the call is the same and only ad server reporting will show that there has been a conversion count. In this call to the server will be carried the extended data parameters as well.

The problem here is that ad server reporting processes conversions slower than the processing of impressions and clicks, so the conversion activity may not display in ad server reporting

for up to twenty four hours. That is a long time to wait to find out that the conversion tag is not firing correctly and that no conversion event was recorded in the ad server reporting. For this reason it is recommended that the creation and implementation of the conversion tag happen at least a few days before the ad campaign is due to go live. If there is a problem with the implementation, it can be tested, remedied and retested before the campaign goes live.

Although it is recommended that the trafficker checks the conversion tag is implemented correctly before the campaign goes live, checking for the "conversion" in ad server reporting for testing purposes has two problems for the tester:

1. The most relevant conversion tags (usually sales tags) will predominantly be housed on a page at the end of the Advertiser site's user journey (also called a "booking pathway"). If this is indeed a thank-you page, the tester will normally be required to undertake a transaction to get there. It is recommended that Advertiser's therefore have 'testing' transactional information that can be inputted into a form (such as a fake credit card number) so that the tester can reach the sales tag without incurring a cost. A surprising number of Advertisers have not foreseen this requirement and as such conversion tag testers (in support teams on the Agency side or working for the third party ad serving provider) undergo personal expense to conduct testing with reimbursement later on. This can result in an embarrassing situation for any Advertiser whose site management team have implemented the conversion tag incorrectly on the thank-you page. Expensive transactions such as holidays can be particularly tough on the tester.
2. If the conversion was actually measured then the reporting will contain the record of a conversion- however this conversion would not really have been a true sale but

a test sale. For this reason some ad servers have conversion tag test functionality built in. This allows the conversion test data to either be quarantined from the actual conversion data for when the campaign goes live or allows the tester to omit their IP address from the collection of campaign-relevant conversion data, so that it only appears in test reporting.

In reporting, a drop in conversions most commonly occurs because the web development team of the Advertiser site are unaware of the significance of the code. A weekend release cycle of the Advertiser site could easily accidentally remove the conversion tag and on a Monday morning, when the development team realise the tag has been removed, they would be alerted in ad server conversion reporting and need to replace the tag. The trafficker would need to just export the same conversion tag code again to fix this. It is wise for the trafficker to check conversion reporting in the ad server reporting interface frequently to ensure that the conversion tag has not been removed. Unfortunately despite programs designed to educate the site management teams about the importance of conversion tags, such technical mishaps will still occur.

A conversion tag that is setup and tested as working can be used to support the tracking of conversion events for a whole campaign but there are additional conversion tag settings and functionality that are utilised in almost every campaign utilising conversions. These settings fall into three groups:

- a. Piggybacking
- b. Conversion Windows
- c. Attribution Modelling

The remainder of this section will look at these settings and explain how they work, how the trafficker can set and adjust them and for what purpose.

a. **Conversion Tag Functionality - Piggybacking**

Container Tags

Conversion tags (spotlight and action tags) can be setup to trigger the loading of other conversion tags. This trigger usually 'fires' when the page with the conversion tag on it is loaded. Conversion tags with this functionality are called "Container tags" or "Floodlight tags" and the vast majority of conversion tags have the ability to perform this function. The container tag can hold a large number of other tags from a variety of sources, it can also be used to carry third party code. If these additional tags originate from another technology provider: these tags are referred to as 4th party tags.

What other conversion tags or third party code would an Advertiser want to "piggyback" onto a conversion tag? Given the usefulness of conversion tags it is not just third party ad servers that want to measure conversions as they happen. Publishers measuring CPA might want to put a small piece of code in the container tag so they know straight away when sales are being made. In addition to this, any technology that drops a cookie on a user's machine at some time, will want to create a connection between that event and the arrival of that user to this key point in the purchase process: the confirmation page conversion tag.

It is also worth bearing in mind that sometimes Advertisers are using more than one tracking technology (double ad serving) and they therefore want the non-dominant ad server's conversion tag to be piggybacked into the container. The aim of the container is to reduce clutter on the page behind the scenes.

When an Advertiser or Agency completes a trafficking sheet for the trafficker at the start of a campaign, it is customary to supply the trafficker with instructions regarding piggybacking. The additional conversion tags to piggyback into the container tag are sent across as code and the trafficker inserts this code into the piggybacking screen in the third party ad serving interface. Some ad servers provide the facility to label each piggybacked tag with a name to identify its origin or purpose. This way when the trafficker makes changes to the container tag, it is clear which tag does what. This is particularly useful when there are multiple tags for different purposes by originating from the same technology provider (therefore having virtually the same code syntax and being difficult to uniquely identify).

Certified conversion tag providers

Some ad servers such as *Atlas*, also employed a process whereby all piggybacked tags needed to belong to recognised and certified providers. To pass certification, the technology providers of the additional conversion tags would need to sign legal documentation to verify that the piggybacked tags would not be undertaking bad industry practices (such as the collection of PII or data collection which could breach Publisher terms such as Publisher data leakage). This process of certification gets a mixed reaction across the industry as some traffickers find that it significantly slows down the process of piggybacking. In some instances technology providers refuse to sign legal documentation which can mean changing plans to use that provider or the ad server altogether for that campaign.

Piggybacking essentially means being able to adjust technology sitting on the Advertiser's site without disrupting the site management team (which will almost always slow down the process of tag implementation). This level of accessibility usually means that swapping tags in and out of

container tags becomes a daily activity for the trafficker, sitting completely outside of campaign launch dates.

Which conversion tag fires first?

When a container tag contains several conversion tags, the container will call the third party ad server first, then the piggybacked tags fire off using a random algorithm. This way the piggybacked tags are all treated fairly and priority is not given to a particular one. Sometimes the trafficker may want to over-ride this default which is where selective pingging comes in.

Selective Pingging

On a high traffic web page, such as the confirmation page of a major Advertiser with an ecommerce site, the container tag might house twenty or more third party conversion tags or tracking pixels (the phrase 'pixels' is used because it is typically the `<noscript>` code in the generated conversion tag code, which is implemented into the container). Calling all of these pixels every time the confirmation page loads, risks the user leaving the page before all pixels have had a chance to load. If the user abandons the page before the container fully loaded then some of those third parties won't track anything at all.

The solution to this is to have particular rules in place on the container tag (controlled from within the user interface of the third party ad server) which have piggybacked conversion tags fire off at different times. The most common method is to limit the firing of those tags which had nothing to do with the sale. For instance, the ad server can check in a matter of micro seconds if the purchasing user cookieID was referred to the Advertiser site from an ad on a specific Publisher site; it then fires the corresponding pixel and ignores the other Publisher-supplied pixels. This method is sometimes called "Selective Pingging".




Figure 4.1.5 In selective pingback the trafficker has received several conversion tags or pixel code from the Advertiser or Agency who have explained that the code comes from a couple of the suppliers on the media plan. These conversion tags (referred to in this diagram as 4th party tags) are loaded by the trafficker manually into the ad server piggybacking interface for the confirmation page sales/container tag. The trafficker then adjusts the selective pingback logic so that if the user arrives to the confirmation page that was originally directed to the Advertiser site by an ad on Publisher A, the 4th party A tag would fire and the 4th party B tag would not.

Piggybacking with Extended data

Piggybacked pixels can capture information that the container tag already collects such as OrderID or any metric and string extended data. To pass the information through,

the ad server utilises an "ad server macro" which is a tiny piece of code, already defined by the ad server which can be placed into piggybacked tags just as a token would be inserted into an impression or click tracker. An example of such a macro would be %orderId% which would be appended to the end of the 4th party pixel code to automatically recognise the Order ID in the container tag and populate the token space containing this macro with the same identifier.

Beyond the extended data, the 4th party tag can also capture pre-defined ad server macros such as the Conversion tag ID. It is worth the trafficker checking with the third party ad serving provider for a full list of the available pre-defined ad server macros.




Figure 4.1.6 The extended data present on the confirmation page is passed into the sales conversion tag by the

customised code on the page (written by the Advertiser site management team). When the sales tag fires, it carries the conversion information with the extended data through to ad server reporting. The sales conversion tag is operating as a container tag and fires the correct 4th party tag using the Selective Pinging logic. Upon firing the 4th party tag calls the 4th party reporting server, carrying with it the same extended data variables present in the sales tag (namely the OrderID and the Revenue).

The conversion tags described here as being piggybacked in the container and the container tag itself will report data back to the ad server when it is fired but ad server reporting will treat this decide whether or not the activity counted as a conversion event using the conversion window settings.

b. Conversion Tag Functionality - Conversion Windows

Window Settings

A conversion tag sends information to ad server reporting each time it is fired, but the ad server has to decide if this information constitutes a "Conversion" or not. We have seen how the ad server makes a decision between a Conversion and an Action but consider this scenario:

A user sees an ad for product belonging to an Advertiser and then switches off their PC, goes on a long round-the-world trip and then returns to the machine two months later. On their return, the user switches on their PC, opens the browser and types in the URL of the Advertiser website before purchasing the same product and firing the conversion tag on the confirmation page after the purchase.

Is this a conversion? Some readers would say no. Some would say that only if the user clicks on the ad, did they

intend to purchase. But in that instance would a two month gap between that click and the final purchase still mean we can attribute the conversion to that click? The answer lies in the ability of the trafficker to customise the conversion tag's conversion window settings.

Post-impression and Post-click windows

There is no "right method" for measuring conversions; but the two schools of conversion measurement are called "Post Impression" and "Post Click" conversion.

- A post impression conversion is when a user has seen an ad but not used that ad as a route to visit the Advertiser site to trigger the conversion tag (the user may have typed the URL into the browser directly to visit the site). The conversion is then classed as a post-impression conversion.
- A post click conversion is when a user sees the ad and clicks on it to arrive at the Advertiser site. The user then triggers the conversion tag from their by visiting the page where the conversion tag is placed.

These two models tend to fit different product types and purchase cycles. Consider the user exposure to a high impact video ad for a vacuum cleaner. The ad server reporting records just one impression to that ad by a unique cookieID. The next event recorded by the same cookieID is a purchase of that vacuum cleaner on the Advertiser site just one week after the exposure to the ad. If no other data exists for the same cookieID the Advertiser can infer that the user converted as a result of only seeing the ad and this is recorded as a post impression conversion.

The time-lag of one week is a consideration that needs to be taken into account based on the fact that the product in question is a high value purchase (a vacuum cleaner) and

that the user may have taken some time to come to a decision to buy the item after seeing the ad.

The "Post Impression window" can be controlled and set by the trafficker on the conversion tag sitting on the confirmation page of the Advertiser site. Given that vacuum cleaners have longer purchase considerations than smaller retail goods (say a week or more), the post-impression window should be set to be over 7 days. The trafficker goes into the conversion tag settings inside the third party ad server interface and adds this setting before the conversion tag is implemented. Here, the post impression window should be set to a value over 7 days. In this example let's suggest that the trafficker sets the window at 12 days.

This means that when the ad server reporting receives the conversion event from the fired conversion tag, the ad server would count this user as having made a post impression conversion. If the user had made the purchase on the 13th day after seeing the video ad, the ad server would not count this as a conversion at all.

Post-click conversion windows work in the same way but consider the time lag between the click on the ad and the conversion activity by the same user. Most ad servers default the values of these windows at 30 days for post impression and 90 days for post click. It is also the case that as a default, the clicking of an ad is a more valuable metric than the impression count on an ad by the same user. If a user's activity for an ad campaign leading up to a conversion consists of a click in the first week and an impression in the second (for a conversion window of 3 weeks), then the click will be attributed as the main driver to the user making the final purchase in ad server reporting, once the conversion tag is fired.

Long conversion windows and short conversion windows

Long conversion windows give the user plenty of time to make a consideration before converting but the question is: how long is too long?

Typically if a trafficker is too liberal and leaves a conversion window open for too long, then user's conversions will be attributed to an ad that may have had no real-world influence on the user's decision to purchase/convert. Here, a user may scroll past a DCO ad and hardly notice it on day 1 and then see a more engaging TV ad on day 20 and convert on day 21. In ad server reporting the attribution for the conversion would be awarded to the DCO ad. Perhaps if the conversion window was shortened to 5 days (so that the TV campaign would not interfere with the online campaign) the DCO ads could be dropped from the campaign since they would appear to provide no direct influence.

However, leaving the window too short means that the true value of the ad campaign may not be correctly attributed. This may be because some users were been influenced by a specific rich media ad but delayed their purchase for too long, meaning that the influence of that ad was not captured in ad server reporting. When the trafficker looks in reporting and sees no conversions for the Rich Media ad (because the conversion window settings were too short) and the trafficker decides to stop serving the rich media ad, the total number of sales may drop.

Ad server reporting can also provide good reports for "time to convert" such reporting shows when most users are converting after seeing or clicking on ads in the same campaign as the conversion tag is associated with (see

Chapter 5). In such reports, considering the collected extended data can help the Advertiser to determine what the conversion window settings should be. Using this method, when the purchases of the product featured in the ad campaign are being made, using a string such as 'Product Name' the Advertiser can spot the difference between conversions for the product featured in the ad from those that were not featured. Those products featured in the ad are more likely to have been purchased as a result of the ads they feature in, so long as the converting users saw those ads in the first place.

Given that there are no hard and fast rules to set these windows; experimentation is key to finding a happy medium, otherwise too many or too few conversions could be awarded to the wrong ads or attributed to the wrong Publishers when paying out on a CPA.

Shadow Tagging

Some ad servers allow the Advertiser or Agency to view their total conversions in reporting by more than one conversion window so that stats can be compared (answering the question; is there a difference in the total number of conversions between short and long windows?). Those ad servers that offer a fixed window on the conversion tag with no option to make changes in reporting need to utilise the functionality of the container tag to see what the effect would be of adjusting the conversion window. In this instance multiple conversion tags are created under the same Advertiser with different window settings. These conversion tags are then piggybacked onto the container. Comparing the conversions across the different conversion tags in ad server reporting exposes the effectiveness of the different windows. This method is sometimes referred to as "Shadow tagging".

Reporting on Conversions

The presence of conversion windows means the ad server has to do a lot of work to calculate conversions as it must look back through the collected data for the entire length of all windows in order to allow reports to be pulled with an accurate reflection of total conversions. The data processing for ad servers which typically process billions of impressions each week means that accurate conversion data (at time of publication) is not currently available in real-time unless the ad serving technology has considerably less information to process. Data is typically available within twenty four hours of collection for reporting purposes owing, in-part, to the processing of conversion window information.

Repeat Conversion Windows

Sometimes users will purchase something and then refresh the confirmation page. To prevent the ad server from counting more than one conversion event, there are adjustable settings in place called the "repeat conversion window" which is usually set at about a minute or more. This way any refreshes of the conversion tag which take place by the same cookie in the same minute, are ignored in conversion processing and are not counted as conversions.

c. Conversion Tag Functionality - Attribution Modelling

It would be unfair to award the conversion to the last ad that a user clicked on, since this would suggest that this last ad was completely responsible for a purchase or conversion that followed it. As consumers we usually experience a number of engagements with a brand before a purchase. These engagements or "touch-points" (a reference used by *Flashtalking*) should each be given some credit for the final conversion. In practical terms this means that rather than

paying the Publisher responsible for the last ad that was clicked, (as would occur under an agreement to pay the Publisher for CPA) all Publishers that contributed should be awarded a portion of the CPA. The division of that amount of capital (even at time of this book's publication) very rarely lands in the lap of more than one Publisher, owing largely to the complex payment terms that would need to be met by all parties. This includes the lack of technology to support such payment systems, and a lack of universal industry agreement in any attribution model beyond what is called the "last click" attribution model.




Figure 4.1.7 If attribution modelling did not exist then only the last Ad touch-point on the user's journey through to conversion is responsible for the purchase/conversion. Here, the user saw three ads for the same campaign across three different Publishers before clicking through to the Advertiser site where a conversion was counted. In ad server reporting, under the "last click" attribution model, only Publisher 3 deserves the full CPA. As a default most ad servers count conversions under the last click model unless adjusted by the trafficker. Using the model can be adjusted under the conversion tag settings screen in the third party ad server interface.

Attribution Settings

Attribution settings present themselves in third party ad servers as a way to adjust the reporting of the total number of conversions as they appears in ad server reporting when looking at conversions driven by each of the Publishers and channels on the media plan. These settings are typically simple adjustments from the default last click model, to perhaps a model that weights impressions and clicks equally and will therefore award the conversion to the last Publisher (even if that Publisher did not have the user directly click on that ad, but rather the user just saw it - counting a post-impression conversion).

Adjustments to these core attribution settings are useful to reveal the true number of total conversions per Publisher and per media channel, when a last impression methodology does not seem to award enough conversions to a more costly channel.

For instance consider this scenario:

- The trafficker begins by setting up the ad campaign according to the media plan. The conversion tag is generated with a "last event" (click or impression) attribution model selected.
- During an ad campaign for 'ski boots' Paid search was the most expensive channel used by the advertiser.
- Alongside this the Advertiser has retargeting tags on their site and is running DCO ads displaying creative chosen based on retargeting. The DCO display inventory is costing the Advertiser less money to run.

- The user starts by actively searching for 'ski boots', sees a Paid search ad and navigates to the Advertiser site (the ski boots product page).
- The user then leaves the site and considers buying the boots but defers by a day or so.
- The user then checks their email and the DCO ad is served to them featuring the ski boots, the product appears to have followed the user to their email inbox.
- Rather than click on the DCO ad, the user navigates directly to the site to purchase the boots.
- The user buys the boots, on the thank-you page the conversion tag fires to signal a conversion.
- In ad server reporting the total number of conversions on the conversion tag attributed to the DCO ad is 1 and the total number of conversions attributed to the search ad is 0.

If the trafficker had selected a "last click" attribution model, then Paid Search, being the last click before the conversion event, would be awarded with the conversion. Here, the attribution settings can be used to better represent the value that search is driving, awarding the whole conversion to the Publisher and channel responsible for the last click.

If in the same scenario, the user had actually clicked the DCO ad as well, a more complex attribution model would have been required to tell ad server reporting that Paid Search (regardless if it was the last click or not) is always more valuable, perhaps DCO because Paid Search better represents the users intent to purchase.

Engagement Mapping

Some ad servers present more complex attribution settings than just "last click" or "last event". Like the more simple settings, these can be adjusted inside the settings on the conversion tag by the trafficker. Rather than make changes to the "conversions" column in ad server reporting, an additional column is presented alongside the original conversion column. Inside the third party ad server *Atlas*, the trafficker can observe the "e-conversions" column alongside the standard conversions column so that there is a yard stick to understand the context of measuring conversions under new models. E-conversions consider how all Publishers, ads and channels can share part of the whole conversion so that the split conversion, in its parts, are properly attributed to a combination of Publisher, Ad and Channel. In this way Publishers could appear to get 0.2 of a conversion in the E-conversions column in ad server reporting. This is determined by the value of the factors involved in true conversion attribution.

As a result of this information the planner or buyer can better calculate the value of the media space back to CPA since a CPA reflects the value of the attributing media space that drove the conversion event. Where a Publisher agrees to be paid on a CPA, the contractual obligation to pay out on conversions should be tied back to an agreement about the attribution model associated with that conversion. This encourages an industry based around a "last click wins" model, to a model more aligned with how users choose to engage with Advertising and subsequently spend with the Advertiser.




Figure 4.1.8 On the user's pathway to conversion, there might be an engagement with touch points across multiple channels. Once the user arrives on the Advertiser site and triggers a conversion tag, the conversion event is fed through into ad server reporting for standard conversion calculation. Within reporting one conversion is counted. In addition the conversion might be processed by a layer sitting on top of the reporting interface known as Attribution Modelling. Shown in the diagram the trafficker has set the attribution model for this specific conversion tag as a 'balanced model' meaning that the attribution is evenly split across all publishers and channels. Inside ad server reporting the three touch points receive equal weighting and e-conversions are a percentage of the overall conversion. Here shown at 0.3 conversions for each touch point. Under this model reporting would show that the Display channel is more effective at driving conversions than Search since the total for Display will be 0.6 where Search will stand at 0.3. Without attribution modelling the final touch point was Display which under the 'last click' model means full attribution goes to Display. With this as a reporting methodology, the correct action would be to take Search off of the media plan for future activity. Such action may easily

be damaging to a campaign since search may be playing an important part of this pathway to conversion, an attribution which is hidden without Attribution modelling being enabled.

Custom Attribution Modelling

Trying to determine which ads are driving actual conversion events is a complex task. On the average media plan, there is a diverse range of Publishers across multiple channels showing different pieces of creative. If a user encounters each of these types on their "pathway to conversion", where does the attribution lie?

In science and mathematics, breaking down a problem into its simplest parts makes it easier to understand. In third party ad serving, particularly inside the Atlas ad server for attribution modelling these parts are called "Factors".

Possible factors include:

- The dimensions of the ad
- The format of the ad
- How recently the ad was encountered (termed "recency")
- The quality of the Publisher delivery (see Chapter 5 - Verification)
- The quality of the creative
- The viewability of the ad (see Chapter 5- Verification)
- The Publisher or site name

An advertiser can rank each of these factors in terms of prioritisation in how to determine the value that each Publisher, Ad and Channel is bringing to a campaign. This prioritisation produces an algorithm to calculate what Atlas terms "e-conversions".

These algorithms are heavily scrutinized as being far too subjective, because the Advertiser has the freedom to tweak

the algorithm in favour of certain Publishers or formats. This can unfairly weight a campaign based on Advertiser and marketing bias rather than a closer examination of the historical cookie level data.

Seasonal fluctuations in the data should also be taken into account, and for some attribution service providers, machine learning is introduced to make light work of building a complex and dynamic attribution model. Given the complexity of such tasks, various point solutions have cropped up to provide such services which are commonly integrated with third party ad serving technology by receiving a feed of their data or connecting via an API. Examples include *Encore Metrics* and *ClearSaleing*.

Many Agencies and Advertisers prefer to make their own minds up with their Publishers about what attribution modelling means for them and they will typically use the data from the Exposure to Conversion (E2C) or Pathway to Conversion (P2C) reports to do so (See Chapter 5 - Cookie level data). In this reporting, all touch-points can be exposed and the Advertiser can look at the whole data-set to decide where the value is being provided on a case by case basis. Large Agencies sometimes have in-house systems to calculation attribution more fairly.

Challenges in implementing conversion tags

Advertisers find conversion tags exceedingly useful and become enthusiastic about ad server reporting when conversion counting is added to the mix of reporting metrics but there are two hurdles to overcome to make the most of conversion tags:

1. Conversion tags take the site management team a long time to implement and a long time to remove. Sometimes it can take a marketing team six months or more to get new tags placed on an Advertiser site.

2. There are so many pages on the Advertiser site: because each conversion tag is different, coding up a large site can be a nightmare if conversion tags and container tags are created one by one

These two needs alone birthed a new frontier in conversion tag management called "Universal tagging".

Tag Management Systems

To get a complete view of the user journey and get a proper handle on conversion attribution, conversion tags should be placed on as many pages of the Advertiser site as possible. This means a lot of work for the Advertiser website managers who keep dropping these pieces of code, received from the trafficker onto pages. Keeping a log of conversion tags across the site can be a full time job in itself. Although container tags have reduced the work by making individual pages a little tidier (keeping all conversion tags inside one container), a large Advertiser site may have thousands of pages.

Tag Management Systems are single pieces of code that (like site side analytics tools) can be included on all pages of the Advertiser site making the implementation very easy to scale. The site management team can simply deploy the single piece of code using the CMS (which will usually grant the ability to update all pages at once). This single act allows the Advertiser to map any existing conversion tags, new conversion tags, retargeting tags and third party code to existing URLs from inside the ad server interface. The role falls to the trafficker to generate and implement the conversion tags all from within the same system.




Figure 4.1.9 Universal Tagging is enabled through a Tag Management system (which can either sit within the ad server or be supplied by a point solution). The trafficker accesses the tag management system in the first instance to deploy the universal tag code which the Advertiser site management team deploy across all pages of the site. Now via the Tag Management system, the trafficker can map existing conversion tags and container tags to pages or URLs on the Advertiser site without needing to bother the site management team to make changes.

Universal Tagging Security

Universal tagging allows marketers a 'backdoor' into web pages on their own site without needing to wait for the website release cycle to come round. With such power comes the responsibility to ensure that any implemented tags and code are tested in full before deployment. Some such solutions also offer the ability to issue a proposal to the site management team for them to grant permission to implement tags via the ad server on a tag by tag basis. This is a security mechanism that prevents the trafficker from inundating the pages of the site with heavy tags or code that is outside the policies adopted by the site management team.

Tag Management Systems are sometimes packaged to include features that measure the speed at which the pages load, which can be an important factor in determining a user's interest in the Advertiser's pages beyond simply firing a conversion tag.

Typically an Advertiser or Agency will deploy such a system long before campaigns are due to go live and then conversion tags will be mapped to the relevant URLs closer to the ad campaign launch date or during the campaign lifetime.

Once again "point solutions" have cropped up in the industry to provide this service at an additional cost to the Advertiser, although Universal tagging solutions built into the ad server tend to be cheaper but with fewer features. Examples of such point solutions include *Tagman*, *Qubit* and the server-side tagging solution *BrightTag*, all of which can be somewhat integrated into the available third party ad servers using piggybacking. Beyond this Universal tagging is very useful to the world of first party data collection and the DMP (see Chapter 8).

Conversion tags provide much needed accountability in digital advertising. Advertising dollars can be attributed to events that provide a direct return on the investment in the ad campaign. Conversion tags allow marketers to learn more about their customers, transforming the behaviour of users into the definitions of audiences. Audience identification allows marketers to correctly target relevant ad campaigns to users and avoid wasting valuable media spend on showing ads to users that will not be interested in an Advertiser's goods or services. Conversion tags do this by allowing the marketers to correlate conversion data across all users that have been engaged in an ad campaign

and check that target audiences are the same users making the purchases.

The next section looks in depth at audiences and targeting to transform the delivery of ads from mass reach to relevance in order to drive valuable conversions.

Chapter 4 - Section 2

Audiences and targeting

- Audiences and Targeting at a media/inventory level
- Audiences and Targeting at a creative level
- third party ad server targeting rules
- Layered Targeting
 - *Figure 4.2.1*
 - Geo-targeting
 - Behavioural targeting
 - *Figure 4.2.2*
 - Contextual targeting (page-scraping)
 - *Figure 4.2.3*
 - Search keyword targeting
 - Targeting based on Verification Data
 - Additional Data targeting
- A/B Testing
- *Figure 4.2.4*
- Tru lift
- *Figure 4.2.5*
- Audience Assignment
 - Channels which bend the rules

Setting up a campaign for any channel within the third party ad server can be a tricky undertaking for a trafficker new to the process, but practice makes perfect. Beyond assigning the ads to placements and sending tags to Publishers, the trafficker has the opportunity to take advantage of more advanced settings in the ad serving interface to ensure better performance (improved interaction, click and conversion rates) from ads in a campaign.

Audiences and targeting settings on a placement level, within the third party ad server interface is a different, more subtle approach to improving relevance than make such adaptations in the planning

and buying phase of the creation of the media plan. These two distinctions can sometimes get confusing:

Audiences and Targeting at a media/inventory level

Part of the role of the media planner is to locate an Advertiser's target audience for an ad campaign among the available Publishers and channels in the online space. Choosing the right audience and procuring that inventory is the first step in ensuring that ads are seen by the relevant users.

When Publishers sell 'audiences' to planners, they are selling ad space which they believe attracts a certain demographic of individual. The methodology behind the demographics should be questioned in depth by the Advertiser before purchasing the ad space. To check that the audience being sold is marketer's target audience, planners can check independent audience verification tools like *Neilsen* or *Comscore* before the campaign goes live, or verification reporting during the campaign (see Chapter 5).

Sometimes Publishers define the audience traffic themselves by allowing a sample size of logged in users (with completed demographic information on the site) visiting those pages where the ads are shown to be scaled up based on the volume of impressions to those pages. Publishers package these audiences for sale in the premium space and label it up for sale directly to planners and buyers or by packaging it for sale on Ad Networks and Ad Exchanges. In programmatic buying the audience can be distinguished using additional audience data or with the help of SSP's (see Chapter 8).

Despite the importance, Audience labelling is not industry standardised so particularly in the Exchange space, it pays to ask questions to be able to verify the audience being purchased. Once the purchase of the inventory has gone

ahead and the labelling of the purchased audience by the Publisher has been accepted, the trafficker can use the ad server as a safety net to ensure that more costly ad formats and creative messaging are shown to the right audience.

Audiences and Targeting at a creative level

An 'audience' in third party ad serving terms means a group of users labelled by trafficker as having a common set of properties (such as users in a particular geography). These groups can be ascertained by either capturing web-page level information in real time or looking back at a user's cookie history and IP data. This level of audience recognition and targeting does not usually omit the delivery of ads to certain users but rather ensures relevant and cost-effective creative is shown to specific audiences.

This section will look in-depth at the ability of the trafficker to apply rules to ad server targeting, to decide which ad to display to the user when the ad server receives a request from a Publisher to display an ad.

Third party ad server targeting rules

As a default setting, the ad server does not encounter any targeting rules on a placement level when displaying an ad other than to show the main ad where-ever possible and fall back to a backup image if the user's browser does not support flash (and falling back to HTML5 for the mobile channel). In Chapter 3 we also looked closely at the ability for the trafficker to set ad controls to rotate ads or shown them in a sequence.

Targeting rules add a new layer to a decision tree that the third party ad server will execute upon, to decide which creative to display to the user. Essentially a series of questions; the rules can be layered up to form complex queries, but there is always an 'other' bucket at the bottom of the decision tree which contains an ad that would be shown if the user did not fit into any of the

targeting rule descriptions. Such an ad would be described as a non-targeted ad. As an example:

- The marketer purchases ad inventory from a Publisher to target only users with a Spanish IP address.
- The trafficker sets up a placement with the targeting criteria to target users with a specific ad for those users based in Spain and sets up a more generic ad in case a user sees the placement outside of Spain.
- The trafficker traffics the placement to the Publisher.
- The Publisher does not honour the targeting criteria for some of the purchased inventory because there are not enough Spanish IP addresses viewing the site to fulfil the agreed number of impressions on the IO. Perhaps even if the site is a Spanish site, Spanish speaking users are still expected to arrive to the site from other locations globally.
- A user from a German IP encounters the placement.
- The ad server receives a request from a user with a German IP for the placement.
- Since the user does not meet the criteria for the 'Spain IP targeted ad' the ad servers deliver the ad in the 'other' bucket setup by the trafficker.
- When the trafficker checks ad server reporting, the trafficker will know that the Publisher did not fulfil the requirement on the IO and may ask for the Publisher to remunerate the cost of the mislaid impression.

In this example it is clear that the third party ad server provides a safety net. Such a feature becomes more relevant if, for example,

there are strict advertising guidelines for advertising medical products in Germany compared to those in Spain. The third party ad server can be used to ensure that the Advertiser does not breach such guidelines even if the Publisher fails to fulfil the targeting criteria on the IO.

Layered Targeting

Given the range of available targeting criteria, it is clear to distinguish how many targeting rules could be set up in the ad server but there needs to be corresponding creative for each defined audience. Complexity can easily become inefficient, if it takes more work to manage the decision tree than the revenue that the targeting generates. If there is no obvious uplift in ROI as a result of the targeting rules in place, it is important try to simplify the decision tree. Branches which attract larger numbers of impressions, conversions, clicks or interactions are worth keeping but those with just one or two conversions or limited engaging behaviours should be scrapped. Complex targeting requires a lot of creative which can be costly and even DCO (see Chapter 3).




Figure 4.2.1 A user encounters a placement on a Publisher page. The placement makes a call to the ad server to decide which ad to display based on targeting rules or criteria setup by the trafficker. Each rule corresponds to the definition of an Audience inside the ad server interface. In this example, if the user fits the criteria for Audience 1 - Ad 1 is shown. If not, the ad server checks to see if the user fits the criteria for Audience 2 in which case Ad 2 is shown. If the user fails to meet the criteria for either of these defined Audiences then the user is shown the 'other' non-targeted ad - Ad 3. If a user's browser cannot support Flash, then the user will not enter the decision tree and instead will be shown the fall-back backup image.

The third party ad server interface allows the trafficker to setup various forms of targeting designed to match the marketer and planner's intended audience criteria. The trafficker can navigate to the placement level within the campaign setup to apply targeting rules to the ads which are assigned to the placements. These targeting forms are:

- Geo-targeting
- Behavioural targeting (also called retargeting)
- Contextual Targeting (also called Semantic targeting)
- Targeting based on Verification Data
- Additional Data targeting

In addition some ad servers offer the chance to implement an additional frequency cap at the audience level. This can act as an override to the rule in place, bringing the user down to the next section of the decision tree pre-maturely. This is useful so that the Advertiser can ensure that the targeted audience does not suffer creative fatigue. This differs from placement level frequency capping (which shows the user a different ad when the cap is met) here, when the cap is met at an Audience level, the user drops down the decision tree and might slot into another rule to be shown a different placement.

Geo-targeting

Before the ad server delivers the ad, it receives the cookieID and the IP address of the user. The IP address can be used to work what geography the user sits within, with a decent degree of accuracy. Upon receiving the IP data, the ad server runs this against a mapping database based on additionally purchased data from IP mapping providers such as *Digital Element*. This will return a geographical location primarily at a country level but then with a greater granularity in some markets when the internet is more heavily used. The US is one of such countries where breakdown to DMA and state level retains a decent accuracy.

The mapping data across the board is not one hundred per cent accurate because not all IP addresses are correctly mapped to the country the user is sitting within. It is sensible to suggest that IP mapping is closer to 80% accurate

but this is still enough to perform geo-targeting, especially if this data is paired with other data sets for targeting. Beyond state or region levels IP data does not fare so well (such as falling to postal code or zip code level) in terms of accuracy so it is recommended to avoid this level of geographical granularity using geo-targeting alone.

To utilise geo-targeting, the trafficker needs to setup placements so that the ads assigned to them are geo-targeting by selecting the specific country or region from the ad server ad-assignment interface. Once the assignment is complete and the campaign goes live, users in the specified geographies will be shown a targeted piece of creative, if not then they are shown a non-targeted ad. Geo-targeting is useful to add relevance to the creative, so a user sitting in Chicago seeing a Rich Media ad from a flights provider such as American Airlines may see a data collection form inside the ad which can calculate the cost of a trip from Chicago to other destinations. Geo-targeting would ensure that the creative is used where Chicago is a pre-populated field in the form. This creative optimization may save the user time when engaging with the ad and prove to increase the total number of conversion events as users navigate to the site to purchase flights.

Although it is possible for the Publisher ad server to perform IP level targeting (thereby geo targeting ads) the control would then sit with the Publisher. The trafficker in the third party ad server interface would need to set up hundreds of tags to provide Publisher-independent granular reporting and would not be able to stop a Publisher serving a high cost rich media unit to the wrong geography.

Behavioural targeting

Behavioural Targeting or Retargeting is now so extensively utilised in digital advertising that it has been given its own Chapter in this book (Chapter 7). To summarise very briefly: users that have performed a particular action on a website, such as making a purchase of a particular product, are added to a retargeting audience or cookie pool, so that specific creative can be targeted to those users in the future. It is wise to add retargeting pixels to the Advertiser site pro-actively so that a corresponding audience and relevant ad can be set up inside the ad server on the placement level.



Figure 4.2.2 Audiences have ads assigned to them. When the ad server receives a request from the Publisher page via the placement, the audience decision logic in the ad server selects the correct audience (such as the correct geography) and serves the corresponding ad.

Contextual targeting (page-scraping)

The page that the ad is served onto provides great contextual information, since the user has navigated to the page for a reason and this is reflected in the content surrounding the ad. A placement can be setup therefore, to gather information about the page itself before choosing which ad to show. The keywords codes or variables present on the page can be used to inform the third party ad server

targeting system; the trafficker only needs to setup an audience and therefore an ad to match with the corresponding keyword or variable.

To go about gathering the information on the page, the placement can either scrape the page automatically or the site management team can write a script to accompany the third party ad tag, passing contextual information into it. The former of these options is the kind of contextual targeting that *Flashtalking* offers, which is a page scraping mechanism which can lift keywords from the page to inform the targeted ad. The latter example is what is referred to in the *DG/Mediamind* ad server as Keyword targeting; here the Publisher passes a variable across such as 'female54years' which may not appear in the visible context of the page but is available as a variable because the user is signed into the Publisher site. Keyword targeting can be useful for demographic targeting where the audience data is being informed by the Publisher's own database.

Gathering semantic information is invaluable when considering the delivery of ads to pages in Belgium. Because Belgium is a tri-lingual society (French, German or Dutch), the semantic of the page informs the placement which language to display the ad in. In this example the trafficker needs to have set up three audiences with corresponding ads in each of the three languages along with a non-targeted ad. The page can pass a variable or keyword across when the placement requests the ad from the ad server to display the ad with the corresponding language.

Contextual targeting can also be utilised in combination with DCO so that the ad passes a value into a dynamic creative which appears to mirror the context of the page in real time.




Figure 4.2.3 Placement A is an ad supports by keyword targeting. A script is developed by the Publisher site management team to pass a specific keyword into the ad request. The keyword is processed by the ad server decision logic and requests an ad from the corresponding audience assignment. In placement B, the placement code contains a scraping mechanism which processes the semantic of the page, passing a value through to the Audience decision logic. Here the same action is taken as with Placement A and the ad request serves the ad for the corresponding audience.

Search keyword targeting

Search Engine owners *Google* and *Microsoft* have the advantage of looking beyond the ad serving data to look at a user's full search history even if they did not click on the paid search ads tracked by the ad server. This search history can be tied to the user's local cookie data or signed in Windows Live or Google accounts on the local machine. The ad servers of these two companies (*DoubleClick* and *Atlas*) can then match the ad serving cookie with the activity of the first party network/search engine cookie offering a wealth of new targetable data to Advertisers using these technologies. Advertiser can setup their ad campaigns to target users based on specific search terms as keywords that they have actively searched on in the past.

As far as other ad servers go; the keyword data is captured in paid search campaigns which are tracked under the search channel by those ad servers and there is an association between those keywords and the cookie. In the future it is likely that these users could be targeted based on paid search click activity although this would probably constitute a form of retargeting rather than be considered true search keyword targeting (since the user has to perform the action of a click for the data to be collected).

Targeting based on Verification Data

Verification data (See Chapter 5) is already being catered for in the world of targeting on most levels:

- Geographical verification is catered for in Geo-targeting
- Brand Safety and Content Categorisation is catered for in Contextual targeting
- Device verification is partly catered for using the ad server fall back mechanism for HTML5

The outliers for Verification data include verification metrics beyond brand safety which determine the quality of the content on the page. Such metrics would include the presence of a large number of ads (called "ad clutter") or the presence of user generated content (UGC) or commenting. Today ad servers are not detecting these occurrences as part of the standard placement delivery (for premium direct inventory buys) but it is likely they will do so in the future. Within the DSP and Exchange spaces this is a different story completely (see Chapter 8).

With regards to targeting for viewability data, the placement upon loading would only be able to gather viewability data in real time, and as such the % surface area seen and the duration of time the user sees the ad for will be subject to constant change. Ad server targeting is not yet sophisticated

enough to contend with targeting ads based on this information. Instead it may be possible for traffickers to use previous data collection regarding average viewability metrics for ads delivered to the same placement as a benchmark to determine the most suitable ad to serve. Once again this level of sophistication has not yet been introduced into targeting because (at time of publication) viewability is still in its infancy. Once again technology is available in the exchange space to buy inventory and target users based on a rudimentary viewability detection method; detecting for ads served above or below the fold of the user's browser.

Additional Data targeting

All of the targeting methods listed so far in this section are actually using additional data to make targeting decisions:

- Search targeting - Search Engine Data
- Verification targeting - Verification Data
- Contextual targeting - Page Level Data
- Keyword Targeting - Publisher-supplied Data
- Behavioural targeting - Retargeting Data in the ad server retargeting pool
- Geographic targeting - Addition Data set from IP mapping provider

There are a significant number of other data points that could mean that Advertisers are targeting users with more relevant creative messaging and more impactful ads. Unfortunately as yet many of these capabilities have not been built directly into the ad server and require a complex customisation of bolt-on or partner technologies such as DMPs or Data Providers (See Chapter 8). These data providers can sit halfway between the placement and the ad server to nourish the placement call to the ad server for the ad with more information about the audience. In this

instance both Retargeting and Keyword targeting can be adjusted to build out audiences in the ad server from external sources of data. This more complex capability is explored in more detail in Chapter 8.

Beyond external sources of data, the ad server is collect high volumes of data which can be transformed into addition targeting options. These are standardised in some ad servers while others may need to set up custom scripts to perform them:

- Targeting by the age of the user's cookie
- Targeting by the date the cookie was dropped
- Targeting users based on their historical cookie level data
- Predictive Behavioural targeting (comparing the actions and engagement of converting users to non-converting users to make targeting decisions)
- Targeting users based on over-zealous frequency

Targeting methods in a decision tree-type workflow allows a trafficker to use trial and error to narrow down the audience left seeing the non-targeted ad. This methodology is highly useful for targeting users that perhaps are accessing the internet from a non-US IP address by customising the decision tree and making full use of the non-targeted ad. In this way, traffickers can calculate what impact certain creative messaging is having as well as creating various "experiments" to split out the users seeing ads to determine if one ad is better than another. Two examples of these are "A/B testing" and "Tru Lift" studies.

A/B Testing

A/B testing means splitting the users that see a placement out so that different users see different ads and then the results between the two groups are compared. This enables the segmenting of cookies into random test or control groups so that the control group always sees the same ad while the test group sees an ad

that changes. The resulting increase in effectiveness of the new ads can be attributed to the creative displayed rather than unaccountable trends.

A/B testing can be conducted in two ways:

1. The first time the user encounters the placement, they are forced into a group (either the control or test group). If the user sees the placement again, they will see the ad for the group they were originally forced into.

By creating a control group the trafficker can test various creative formats, creative messages and even post-impression windows against users that have been served the ad in the control group. With this method, if an offer is never shown to a user in the control group, yet they actively search for it and convert, that media in other channels (such as TV or word of mouth) may be attributed to the conversion.

2. Every time the user encounters the placement they encounter the logic from the top of the targeting decision tree.

This way the control group changes so that the division between users is relevant only to that session in time. This way if users in the control group had a propensity not to convert any way, the control group would not always stay the same so those non-converters could be balanced between the control group and the test group.

This is not a true A/B test as individual users commonly visit both experiments. However it controls creative exposure equally and more fairly at a specific point in time, to establish the value of creative messaging.




Figure 4.2.4 An A/B test allows a trafficker to control the delivery of specific message to a sample of the total number of users. When the user first encounters the logic they enter the A/B test, 90% of the users will see the usual Ad B, while 10% will see the experimental message Ad A. When the user encounters another placement from the same campaign they are shown the next usual message in the sequence. Did the presence of Ad A instead of Ad B improve the campaign performance? A trafficker can check reporting to see. If the experiment failed to point to a performance increase or worse; if Ad A had a detrimental effect, at least only a small sample of the total users were ever exposed.

A/B testing helps expose the strength of the creative message and its relationship with Advertiser KPIs such as conversion events. A "tru lift" study is the name given to an experiment using A/B testing that allows the Advertiser to demonstrate that any digital

advertising is having an impact on conversions versus users that are not exposed to the campaign at all.

Tru lift

A tru lift study tries to prove that an exposure to an advertising campaign causes (at the most rudimentary level); an increase in the number of conversions for an Advertiser over a given period of time.

To conduct a tru lift study the trafficker needs to split the audience of the advertising campaign into those that have never seen an ad for the Advertiser and those that have ever been exposed to an ad. Given the limitation of the impression based cookie window it is actually very difficult to create such a clean break in the audience; some audience members may have seen Advertiser ads from previous campaigns and therefore be influenced to purchase with the Advertiser on that basis alone.

In an ideal world if the user has never been exposed to an ad from this campaign or Advertiser before, the trafficker should set up targeting to create a retargeting pool for users that have seen ads previously and allow the control group to see the non-targeted ad. This non-targeted ad would not be for the Advertiser but rather a public service a (such as an anti-smoking ad). If the ad server observes the same rule every time the user encounters one of the placements for the campaign they will never be shown the targeted ad and therefore the group remains controlled. The third party ad server allows audience to be split and weighted so that the first time, 50% of the users enter the control group and 50% are shown the Advertiser's ad. For a tru lift study to work, the rule which determines which ad the user sees, must permanently drop that user into a cookie pool.

One problem with this approach is that to be a water-tight control group, the Advertiser must begin delivering retargeting pixels (see Chapter 7) with their ads from day 1 of their first campaign otherwise some users will be influenced by previous campaigns.

The retargeting cookie window also needs to be left open long enough to allow the tru lift study to be undertaken. The other problem is cookie-deletion, which cannot be avoided where a user who was seeing advertiser ads could end up no longer seeing them as they are placed into a control group.

Across the board, tru lift studies are a poor indicator of the power of online advertising. The studies undertaken are commonly full of holes that cannot be plugged and the results the advertiser was expecting to see rarely come to fruition. This means that a lot of time and money can be wasted under taking such studies with little or no return.




Figure 4.2.5 A tru-lift study. The user arrives at a placement, never having encountered the campaign before. The audience criteria asks if the user is in the retargeting pool for Ad A or Ad B. If the user is in neither audience type then they enter the A/B test. Here the ad server randomly separates out the users equally so that

50% of users are in the control group while 50% of users are in the test group. The corresponding ad includes a retargeting pixel so that the user is labelled as belonging to one of the audience groups. The next time they encounter a placement from the same campaign, the logic kicks in and recognises that the cookie is in a retargeting pool. It shows them the corresponding ad for their audience type. Now the control group can always be shown the same ad or even an ad for a public service while the test group are shown Advertiser ads for the purposes of creative testing. In reports the performance across the two ads for both clicks and conversions can show if the creative message in the test group is showing an improvement in performance over users that may have only seen a public service ad.

Audience Assignment

Targeting, experimentation and audience assignment differs between third party ad servers. At its heart the audience is defined by the trafficker and then that defined audience is paired with a particular ad or ad version.

A placement follows the list of audiences down the tree of targeting decision criteria to determine which ad to show. One would think that it makes sense to always try to target an ad to a user to increase the chances of the ad getting noticed and action being taken but a long list of targeting rules for each placement increases the work required to setup and optimize the campaign. High complexity in setup (such as setting up five or more different targeted audience combinations for each placement) causes the following issues:

1. Much heavier processing required at the ad server when delivering the ad, which could impact the user experience of the ad load time

2. Keeping tabs on the effectiveness of various audiences becomes very complex when only very small numbers of users convert on each targeted ad
3. Committing to so many audiences and granulating the setup so much means having to create many ads to support it (although DCO will do this well the setup and assignment can be lengthy and the optimization will not work if the audiences are so small because of the diminished effective sample size).

These can be avoided by using the same audiences for multiple campaigns. In addition, the recommendation is to reduce targeting and audience complexity with consistent optimization off of the back of well analysed reporting with the ultimate aim to use only a handful of very well targeted and effective audiences. In essence targeting is a science experiment that if pursued intelligently can yield great results.

Channels which bend the rules

Audience targeting within the ad server will not work for pure Paid Search Campaigns (since the ad server is not delivering a creative), this is also the case with Affiliate and Email campaigns. Within the Instream channel, targeting is available but since the video is part of a running stream (where a pre-roll is followed by video content), there is no need for a non-targeted ad. In this instance the ad server receives a request from the placement to serve an audience-specific ad. If the criteria is not fulfilled, the ad server returns no information, meaning that an ad never loads and there is no non-targeted ad and no fall back. In this sense the ad server (at least for Instream) has the power to affect the actual media delivery.

In addition to this, Ad blocking (a feature of verification) could also be considered a feature of targeting, where, if the page does not meet the targeting criteria for verification, a non-targeted ad is displayed. However, in this instance the non-targeted ad is a public service image and its delivery is specially separated out in reporting for the purposes of remuneration (see Chapter 5 - Verification).

This concludes the work of the trafficker in the pre-campaign setup interface in the third party ad server.

The trafficker will return to this interface to set up subsequent campaigns for retargeting (see Chapter 7) as well as adjustments to the existing campaign for the purposes of optimization (see Chapter 6). The next step for the trafficker is to access the ad server reporting interface. This will allow the trafficker to check ad delivery, performance and verification.

Section C

Ad Server Analytics

Section C

Ad Server Analytics

The process of Researching the right Publishers for the media plan, followed by planning and buying the media for the campaigns is followed by the process of trafficking, all for the proof of delivery, verification and performance. Reporting and analytics are the ultimate a feedback mechanism for an Advertiser, that yields the results and insights to ensure bountiful effectiveness and future campaign optimization.

The next section looks in depth at reporting, analytics and metrics. Just knowing what data can be returned from ad serving systems, can be enough to find the great in any campaign or improve a failing one. Surprisingly, statistics are rarely difficult to digest in spread sheet form, where the uniform layout of the data in a standardised to read the data quickly, becomes a task easily achieved.

The stories of the data spring out of reports pulled in the right way and at the right time. The ad serving providers continually improving the look and feel of the interfaces and data outputs that make it easier to digest that data and find actionable insights. Even at time of publication, ad server analytics suites are undergoing transformations of this nature. The current state of play comes down to three core areas: ad Verification, ad Delivery and ad Performance, this following chapter will dive deeply into each of these areas in depth but first it is important to understand What reporting is made up of, Where the reporting is generated, When the data is processed and how the data is outputted.

Chapter 5

Reporting Technology

Chapter 5

Reporting Technology

- What is ad server reporting?
- Aggregated Data
- *Figure 5.0.1*
- *Figure 5.0.2*
- Reporting API
 - *Figure 5.0.3*
- When is the Data Processed?
- Real time reporting
- Processing Clock
- Where is the data processed?
- How is the data filtered?
- Where is the data displayed?

This section is an introduction to reporting technology within the third party ad server. It explains what ad sever reporting is made up of. When and where the data is processed and what the outcome of the processing is.

What is ad server reporting?

Third party ad server reporting is an interface that sits on top of the data which has been collected during a live campaign. In the core databases of the ad server, the data is virtually incomprehensible, it needs to be collated and presented in a fashion that is easy to understand, the outcome are the available ad server reports. This section will explore the relationship between the collected data and the reporting interface. This relationship carries certain conditions which put the reports into a context which must be taken into account when analysing the reports themselves. These conditions include how the data is processed, where it is processed and when it is processed.

Conditions are commonly presented alongside fields inside ad server reporting and give clues as to why certain trends may exist in the data.

Third party ad server reporting is divided into two distinct types: Aggregated data and Cookie Level data.

Aggregated Data

Aggregated data describes data that is presented to summarise the count of certain metrics (such as impressions and clicks). Summaries of this data include time periods and ad server hierarchy entities (such as placements and Publishers). These summaries are commonly selectable as filters in the ad server reporting interface such that they control the amount of data that will appear in the generated report. The filtered list of entities will usually be presented in spread-sheet form down the far left column to describe the content of each row. The titles of the columns will be the metric names themselves and the cells will contain a count of the occurrence of each metric against each entity. The design of the spread-sheet is such that a human can analyse the data right across one or more campaigns.

	Impressions	Clicks
Campaign	0	0
Network	0	0
Publisher	0	0
Site	0	0
Package	0	0
Placement	0	0
Ad	0	0
Creative	0	0
Version	0	0

Figure 5.0.1 An aggregated report from a third party ad server aggregates metrics across an entire campaign. The data will be available for each level of the ad server hierarchy. To 'pull' a report of this type required the trafficker to navigate to the analytics interface of the third party ad server and input several filters such as the Campaign name, the specific hierarchy entities and the date range.

The ad server typically takes a few seconds to bring the data together before churning out a spread sheet or visual representation of the request (referred to as a dashboard). Commonly, the ad server has already processed the data it is returning and the delay accounts for the front end of the ad server communicating with the backend storage systems to retrieve the data in the requested view.

Cookie level data by comparison is not available in the ad server interface but is available via a feed mechanism, send very granular data on a given time period to a location outside of the ad server. This information is usually too granular for a human to process and would normally feed into an external reporting system or data warehouse. Cookie level data will be covered off in more detail later in this Chapter.



Figure 5.0.2 Third party ad server reports are available in three main formats. Aggregated reports display high level data which can be pivoted across various hierarchy entities and metrics. A Dashboard will be a more visual representation of the aggregated reports (useful to lift out trends and to present the data). The cookie level data is a more raw output which tends to be a broad spread sheet, table or database in the most granular form available. Here each engagement with an ad or conversion tag is displayed row by row.

Reporting API

In Chapter 3 - Display ad serving, we made reference to the capability of the third party ad server to supply access to reporting data via an API. Different ad servers are able to provide different levels of reporting granularity via their reporting APIs. Almost no ad servers today will allow a developer (with the right level of security access to the API from outside the ad server) to return cookie level data. Only aggregated data is supported by the API. In addition to this not all metrics and ad server entities can be retrieved, so to this extent the API is less detailed than the data that can be collected from pulling reports directly from the interface. Reporting APIs can be of used by developers in conjunction with other APIs provided by the ad server such as a Campaign Setup API or Planning API.

As an example an Agency could hire a developer to build a planning tool for Agency use only which sets up campaigns and placements in the ad server from a distance and colour co-ordinates the buys to correspond with levels of impressions at which ads are being delivered.

The difference between third party ad server analytics and Site Side Analytics

Chapter 2 explained how an Advertiser would install a site side analytics tool (such as *Google Analytics* or *Adobe's Omniture Suite*) onto their website to establish the sources of the traffic coming to the site and to determine how users navigate around the site. Although this data remains very useful, it is still just describing user behaviour on site and cannot explain how user engage with or see ads on the rest of the web. It is true that Site Side Analytics can track referral traffic into the site but attribution beyond the last engagement is challenging unless the third party ad server is integrated with the site side analytics tool.




Figure 5.0.3 Site-side analytics is useful for page-level marketing optimization. The data retrieved from Site Side analytics shows how users arrive to and depart from individual pages and the journeys that are taken when the user is on site. By comparison, third party ad servers will only detect the traffic to a handful of key pages where conversion tags are sitting but reporting will extend out to the plethora of served and tracked ads across channels and Publishers to represent the effectiveness of paid-for advertising.

Having said this, the technologies cannot be perfectly synchronised because the methodologies are fundamentally different. One example is that the ad server definition of a unique user is different from site side analytics (which can adjust its methodology depending on the report time frame. Unique's per month and Unique's per year is not comparable with what the ad server tracks which is one unique user per lifetime of the cookieID.) Site side analytics can also not detect for post impression conversion events and does not offer Cookie level reporting. There are numerous other differences as well around the click methodologies and the position of the code on the

Advertiser page (site side analytics code versus the position of the conversion tag).

There is still value to be attained from Site Side analytics and integrations are possible through the use of token values passed through in the user click pathway. However the two systems cannot replace one another, they are doing two very different jobs.

When is the Data Processed?

The way an ad server collects data and processes it, differs according its core architecture. Ad servers are structured to try to bring data collection as close to reporting as possible in time. While impressions and clicks can be passed into reports soon after user events have occurred, it is not as easy for conversion data, where the ad server must compare conversion events across Publishers to establish the attribution. This means that the conversion information requires a longer processing time. At time of writing conversion metrics are still available after a window of around 6 hours for most ad servers (or 4 times a day). This figure is decreasing as ad server processing power improves year by year to try and reach "real time reporting".

Real time reporting

Real time data allows the Advertiser or Agency to make decisions immediately after users engage with ads. While retargeting makes real-time targeting possible already, if the data needed to be processing by human eyes before decisions were made then the process would slow down. For human eyes to see what is going on, reporting needs to present data, which in itself is subject to a minor delay. His human touch might also be called manual optimization (See Chapter 6) so to speed everything up, decision engines are being created and many already exist, to automatically optimize campaigns as they run to avoid the delay caused by human interaction. Such systems eradicate the need to pull reports for the purposes of optimization.

Having said this, real time reporting remains in demand for direct response Advertisers where a trafficker needs to react to adjust the purchase of media to reflect the data displayed in ad server delivery. For instance: over the course of a day, if a big event happens unexpectedly, the trafficker cannot react to a spike in search traffic for certain terms (so as to pay more or less for certain keywords) since the spike would appear in reporting which may not be available until a few hours later.

If there is ever a fault with the ad serving technology, ad server reporting can usually expect to be delayed while the collected data is correctly presented. In such an instance Advertiser teams can often be waiting to make changes to a media plan; literally prevented from doing their jobs until the conversion data has become available.

Processing Clock

Conversion processing can take several hours, so often the ad server has a hard cut off time for data collection on a particular day or processing period: such as 11.59pm. Then the conversion data will be made available within a few hours, leading up to this cut-off point. The freshness of the data is therefore impacted not only by location of the processing servers but also for the amount of time it takes to process and return the data sets. When pulling reports and comparing data sets for discrepancies it is an important step to ask questions about when the cut-off point was for data processing, what time the reported data was available from. Sometimes an ad server report may be pulled in the CET time-zone but the data it displays was cut-off on an EST time-zone, owing to the location that the data was processed from. This calls into questions where the data is being processed exactly?

Where is the data processed?

Due to the global nature of third party ad serving, the backend systems have a very real and complex consideration to make with regards to what are called "server clocks" (i.e. the local time of the server where the data is being processed). Running all the conversion processing from a single location and then distributed worldwide means the data is available at different times in different countries and really all countries would like the data available as early as possible in the working day, so that Advertisers and Agencies are working to optimize against the freshest data sets.

Matching up data between time zones can be a big headache. Account settings will usually allow the account to be setup for data processing a certain time zone but this isn't always enough a reliable indicator to explain what time periods the report is truly presenting. It is work enquiring with the support team at the third party ad serving provider to see what clock settings are in place and how this impacts the reporting data. In some instances these questions can help to remedy many discrepancies between reports which appear to be for the same time period but were processed for different cut off points.

How is the data filtered?

Data processing does not just mean reporting, it means cleaning up the data. This can mean mapping different tables of data together into a single report, calculating metrics and of course filtering the data. Third party ad servers have their own blacklists for IP addresses that have been compromised and have ways of sweeping the data sets for bots and data deemed to be created by unauthorised machines creating fraudulent clicks or wasting impressions not seen by human eyeballs. Some ad servers allow this information to be presented in reporting. It is also worth asking the ad serving provider if test impressions, clicks and

conversions can be filtered from reporting. Some ad servers have the capability to do this by filtering out activity from certain IP addresses as a backend service. It is also not fair to expect an ad serving provider to divulge the methodology behind the filtering of the data for fraudulent activity since the availability of the methodology might compromise the filters. With this in mind ad servers will usually be industry accredited for methodologies which filter fraudulent traffic.

Where is the data displayed?

Processed data can be accessed via the third party ad server reporting interface. Aggregated reports will exist in templates so can quickly be run (such as a report containing all the standard delivery metrics). The interface will also allow the customisation of the aggregated reports to add specific metrics or filters to the report so that it can be generated by a particular time period or for specific ad server hierarchy entities. Run reports will usually output an excel spread-sheet or in a visual interface where the trafficker or analyst can analyse the data. Aggregated reports can also be shared once created or scheduled (delivery an updated copy of the report directly an inbox on a given time period).

With an overview of aggregated reporting now completed, the remainder of this Chapter will detail out the different kinds of reporting available within the aggregated category before covering off Cookie Level reporting. The sections within this chapter are structured in the order that the trafficker or analyst might retrieve the data from the reporting interface in the ad server once a campaign has been set live. The first stop over is Verification reporting, the first set of aggregated reports that the trafficker will encounter when analysing a live campaign.

Chapter 5 - Section 1

Basic Aggregated Reporting

- QA Reports
- Global reporting
 - *Figure 5.1.1*
- Delivery reports
- Verification Reports
- Performance Reports (Conversion and Attribution Reports)
- Cost and ROI reports
- 404 error reporting
- Expiring cost packages
- Channel Specific Reports
- The Standard Metrics
 - Impressions
 - A new impression?
 - *Figure 5.1.2*
 - Clicks and Click-through Rate (CTR)
 - Conversions & Conversion rate

There are a flurry of standardised aggregated reports that the trafficker can pull from the third party ad server reporting interface. The following section will detail these report types and then examine the most common metrics to be found in said reports.

The basic aggregated reporting types are:

- Global reports
- Delivery reports
- Verification reports

- QA reports
- Cost and ROI reports
- 404 error reports
- Expiring Cost packages
- Performance reports (Conversion and Attribution reports)

QA reports

QA reports are reports that can be pulled before a campaign is put live or from the very first few minutes of a campaign lifetime to monitor key metrics to ensure that the implementation of tags has been conducted correctly. Some ad servers bolster this offering by connecting it to a monitoring platform with a faster data refresh rate than the other data sets. Monitoring platforms such as this can be connected to specific email alerts that can be setup to warn a trafficker when a campaign is delivering or performing unexpectedly. This ensures a trafficker can quickly react to any sudden turns in data trends to perform the relevant optimization (see Chapter 6) or test the live tags for errors. QA reports are particularly useful to check that conversion tags are live and pulling in the correct extended data parameters. Without QA reports, the trafficker would have to wait until data has been refreshed in more complex reporting such as cookie level data, even in cases where extended data collection is just being tested.

Global reporting

Third party ad serving technology can support the needs of international Advertisers and advertising campaigns by offering the single co-ordinating team an access level in ad server reporting that gives a view across all countries. Within this structure, each local country team can only view their in-country reporting data and layers can be added to support regional data access as well while ensuring that only those with global access, see everything.

It is important to remember that when running global campaigns that it is challenging to compare cost data since each country might pay for media in local currency and as such costs need to be standardised to a single currency such as US dollars before being compared with each other.




Figure 5.1.1 Within global reports, there are various access levels to the ad server reporting interface, granted to the analyst or traffickers depending on what level of data they require access to. Centrally managed campaigns (campaigns managed for the whole world from a single location) benefit from this style of reporting, where the supporting teams in local markets see less information in their reports. Access can be granted at Global, Regional and Local levels. Consistent naming conventions are essential to producing "clean data" in global reporting. A consistent global Advertiser, campaign and media plan naming convention (1) allows the global Advertiser to compare performance and

delivery across the regions at a brand and product level. A consistent Publisher name (2) allows for global benchmarking so that the delivery and performance of the same formats can be compared showing if Publishers and Placements are just as strong region to region. A consistent placement, ad and creative naming convention (3) identifies how the creative and execution compares between locales.

Delivery reports

Delivery reports aggregate data about the core ad serving responsibility: i.e. Serving the ads. Such reports can be very useful to Advertisers that have a strict delivery plan, to check that they are on course to deliver the paid-for impressions over the intended time period. Such reports are commonly shared with participating Publishers who can use the data to adjust the controls in a Publisher ad server for over or under delivery. It is best practice for Advertisers and Agencies to check delivery reports at least once per day during a campaign.

Verification reports

The trafficker would run and look at a verification report when a reasonable amount of impressions have been served to produce a large enough data set for the purposes of ad and campaign verification. It is recommended that at least one day passes on longer term campaigns before checking verification reporting. Verification reports are very extensive and a growing area for third party ad serving, as such there is a full section covering this area following this section within Chapter 5.

Performance reports (Conversion and Attribution reports)

Once the ads have begun delivering the trafficker can check in on how the ads are performing to drive certain marketer KPIs. Performance reporting displays the numbers of clicks and conversions broken down by Publisher, Site, Placement

and Ad. For each of these ad server entities, the trafficker can view the total number of clicks on the ad, but then a breakdown or drop-off to each of the conversion tags on the Advertiser site. This shows which creative is driving the most effectiveness. In addition, the e-conversions attribution, explained in Chapter 4, can be displayed alongside the figure for total conversions to each entity. Where it is possible to do so, the metrics from performance reports should be combined with verification reporting so that post impression conversions that occurred on impressions that were not in view, are omitted from the conversion tally.

Cost and ROI reports

By inputting the media cost data into the campaign management interface, reporting can show the granular detail of actual expenditure over time, showing how much it has cost the Advertiser to advertise to different Publishers, Sites and Channels. Furthermore this style of aggregated reporting links performance (conversions, clicks and interactions) back to cost to show how much engagement, exposure and conversion is won from how much media spend. Some ad servers will go a step further and display the ad serving costs alongside the media costs so that the Advertiser can work out the total cost to serve and deliver an ad. When the metrics from cost and ROI reports are combined with the metrics from verification reporting, the Advertiser is informed how much media spend is being wasted on ineffective impressions.

404 error reports

Third party ad servers that do not have a close-to-real-time monitoring and alerting report functionality, offer weekly reports that show which ads are reporting which ads are clicking through to non-existent pages. This is a very useful report to check at the start of a campaign to ensure that ads are clicking through correctly. Those redirects that land

users on non-existent pages through up an HTTP 404 error which can be reported by the ad server.

Expiring cost packages

Some ad servers offer a report which provides another alert-style system, showing when certain cost packages and placements are nearing their end date. Sometimes ad campaigns can last a year or more and although traffickers will insert the campaign end date at a point in the future, these end dates can quickly creep on and if the campaign hits the end date, start serving fall back creative such as backup gifs.

Channel Specific reports

Some third party ad servers group reports together (which are really channel specific reports) into categories like Engagement reports (for Rich Media interaction reporting). As there are an extensive number of channel specific reports, with metrics that need detailed explanation, there is an addition section dedicated to this topic within this chapter.

The Standard Metrics

Outside the channel-specific metrics almost all aggregated reporting comes back to a handful of core metrics. These metrics have already been referenced numerous times in this book, the following is a deep focus on exactly what each means:

Impressions

Most ad server reports, flooded with numbers, represent the increment of impressions for a particular metric. To be found inside the core aggregated reports within the ad server reporting interface is a column purely called "Impressions". Strictly speaking, impressions are a count of the number of ad requests coming from a user's browser for a single ad. The request is counted when it is received (see

Chapter 1) in accordance with the IAB industry approved standard definition of an "Impression". This is what is referred to as the standard industry "trading currency", since the majority of cost models are performed on a CPM (cost per thousand impressions). Ad serving in particular is charged to the Advertiser on this same cost model.

A new impression?

Impressions may have had their day. There are now a few potential candidates in the works that might push impressions off of its held spot as the "trading currency" for digital advertising but there is contention over what might happen. The first contender, which has cropped up a number of times over the short years that digital advertising has been around, is "Delivered Impressions", also referred to as "ad loads" and "ad starts". Rather than count on the ad request, the count is made when the first frame of the ad is loaded. The delay between the standard impression and the delivered impression can be significant enough for Advertiser to question the standard metric imposed by the IAB.

The second contender is the viewable impression (see the next section in this chapter) which only registers an impression when the ad is in view on the users screen. Again there can be a large difference in the number of counted impressions on the ad server versus the number of viewable impressions.




Figure 5.1.2 As the user arrives on the Publisher page (1) a request for the ad is sent to the ad server. The counting server registers an impression as it deploys the impression tracker followed by the ad (2). Ads are usually loaded in this way, so those below the fold of the user's browser will also load up and an impression will be counted (3). If the user scrolls down, any ads which then are shown above the fold will trigger a "viewable impression" metric (5).

Some argue that the industry factors in the difference between the cost of an impression and the cost of a delivered or viewable impression by keeping a low cost where a delivered impression or viewable impression may be provided at a higher charge, because of the delivery verification that surrounds it.

Clicks and Click-through Rate (CTR)

Clicks are a recorded increment in the ad server, counting a user clicking on an ad. As a component of the click redirect, the ad server is pinged and the master database of the ad server appends a row stating that a particular cookieID

clicked on an ad at a particular time from a particular Publisher site.

For many years, the click was the only way a user could interact with an ad and show some intent to purchase. On the early web, Advertisers and Publishers were obsessed with the CPC model. With other methods of counting now available such as rich media interaction or view based conversions, the click has lost a lot of value in the display space. Having said this, a click is some indicator of interest and many Advertisers try to separate out their audiences by those that show some interest from those that don't.

Clicks remain integral to digital advertising owing to the support of them by the performance channels (Affiliate and Paid Search). Clickthrough-rate composes of the number of impressions divided by the number of clicks to determine how many viewers successfully became clickers. Usually this can be expected to be a very low number, certainly much lower than 1% of the total impressions. On a big enough scale this is still a considerable number of interested users, but it is difficult to determine what caused the click itself. Sometimes the creative may have been the most enticing part of the click action (creative optimization) and sometimes it may just been down to the correctly targeting audience (media optimization). As explained in Chapter 4, experimentation can help determine what is driving the user to click.

Conversions & Conversion rate

Just as with clicks, conversions are presented in aggregated reporting as a measure of campaign success, particularly where the user is driven to make a purchase online. Conversions are really; impressions for which a conversion was reported. Conversion reporting will typically displayed so that the total number of conversions is attributed to each

ad server entity (channel, Publisher, site, package, placement, ad, creative and version). Conversions to a particular conversion tag will be labelled in reporting along with a conversion rate (conversions divided by impressions) to calculate how many of the users that were exposed to the ad, went through to convert.

As a reminder from Chapter 4, conversions can be broken out into post-impression and post-click conversion events, allowing the Advertiser to dismiss conversions based on impression only exposure to the user. Furthermore it is important when pulling basic aggregated reports (such as performance reports) to remember that conversions are calculated based on the settings shown on the conversion tag in the campaign management interface of the ad server. Both attribution settings and conversion windows will have an impact on the total number of conversions displayed in the conversions column.

With a good grasp of the basic reports, a trafficker can go into the ad server reporting interface and pull the report directly or setup a schedule so that it delivers to the trafficker's inbox frequently. Third party ad servers design these basic aggregated reports in the interest of time saving so that the trafficker does not have to assemble a report from scratch by collecting up and filtering by a wide range of metrics and hierarchy entities that may not fit together. An example would be trying to report a channel specific metric across a campaign running in a channel that does not support the same metric.

Further in this Chapter we explore Advanced aggregated reports, which the trafficker would attend to a week or so into the campaign, so that enough data has been collected to be able to action the results.

Leaving behind the campaign setup and having the Publisher set the ad tags live means that the work of the trafficker now moves to monitoring and reporting on the progress of the campaign. The first stage in this process involves checking that the campaign has actually kicked off (i.e. are the ads delivering, clicking through and are the users converting). It is recommended therefore for the trafficker to check basic aggregated reporting within twenty four hours of the campaign going live. The initial data available in delivery reports would just be an indicator that things kick off well. The trafficker would then move onto looking at verification reports before taking a deeper dive into the available data via more advanced aggregated reports, fully customised reports or channel specific reports.

The next section looks at this second step: verification reporting.

Chapter 5 - Section 2

Ad Verification & Viewability

- Effective and Ineffective Impressions
- Reactive, Preventative and Pre-Emptive Verification
 - *Figure 5.2.1*
 - Publisher and Placement quality indicators
 - Contextual Verification Quality Indicators
 - Content Classification
 - Brand Safety
 - *Figure 5.2.2*
 - Demographic or Audience Verification
 - The panel-based Audience Verification Method
 - *Figure 5.2.3*
 - The social profile-based Audience Verification Method
 - *Figure 5.2.4*
 - Discoverability Verification Quality Indicators
 - Geo Reporting
 - Device Verification
 - Viewable impressions
 - *Figure 5.2.5*
 - Double Implementations and Fraudulent Serving techniques
 - *Figure 5.2.6*
 - Page Level Verification Quality Indicators
 - Domain URL Reporting
 - Other uses for Page Level information
 - White Lists and Black Lists
 - *Figure 5.2.7*
 - See through, View through and Recordability Rates
 - *Figure 5.2.8*
- Preventative - Ad blocking
- Pre-emptive - Verification in Programmatic
- The Verification landscape

After checking delivery reports briefly to see that the campaign started without a hitch, the next most important area for the trafficker to look into is Verification reporting.

Effective and Ineffective Impressions

Verification is a means of auditing the delivery of the ads. By consulting Verification and delivery reports early and often in campaign management, the wastage of ads can be minimised. Impressions which are not wasted are sometimes referred to as "effective impressions" while the wastage would be "ineffective impressions".

Verification tools have only just started to appear integrated into the third party ad server. Up until 2012 the role of the Verification toolset has been played by smaller independent technology companies but separate tools mean a bigger headache for the setup of the campaign and a larger room for reporting discrepancies. Ad servers with integrated Verification mean that enabling Verification can be as little work for the trafficker as flicking a single switch.

Verification asks a simple question "Is my campaign being delivered to the right audience?" Incredibly this is still a conundrum. It's up to the Publishers to ensure the advert gets to the right eyeballs but being a third party, the ad server allows the Advertiser to verify the delivery independently. Verification tools started life seeing if ads were delivered into a "Brand safe" environment i.e. no adult content, violence or disaster news stories and now offer something wholly more attractive to Advertisers.

Reactive, Preventative and Pre-Emptive Verification

At the time of this book's publication, Verification is one of the hottest topics on the lips of Advertisers, Agencies and Regulators.

Verification can spell trouble for Publishers who defy targeting and audience criteria and are out to make a fast buck by bending the truth about what they are selling to Advertisers. Verification is available to Advertisers and Agencies through three methods:

- **Reactive:**
as in detecting that the ad was served into an unsuitable environment and then reporting the fact in ad server reporting. Such reporting would be used to leverage conversations with Publishers for better targeting in future or remuneration.
- **Preventative:**
as in preserving the delivery of ads and creative messages from unsuitable Publisher pages by serving a public service ad instead. This is currently being called "Ad Blocking".
- **Pre-Emptive:**
as in verifying that an ad will be served on a suitable page and in a suitable spot before the inventory is even purchased (currently only available in the programmatic and RTB spaces - see Chapter 8).




Figure 5.2.1 When the user arrives on the Publisher site, it is possible to verify that both the page and the user are suited to the ad being displayed. An Advertiser has three options. Either the Advertiser buys the ad spot programmatically (see Chapter 8) or by buying directly from the site. In the first instance a pre-emptive tool can be triggered when the placement loads, checking the page content to verify if it is a good place to serve an ad. If the page does not appear to be brand safe, the pre-emptive tool will ensure that the ad spot is never purchased in the first place (sending the available impression back into the Ad exchange). If the Advertiser does not use a pre-emptive tool then Preventative technologies like Ad Blocking can be enabled on the ad server (so that if the page does not look brand safe, the ad server will serve a Public Service Ad (PSA) instead). If neither pre-emptive nor preventative methods are being used, the Advertiser can be reactive and detect the delivery of the ad on an unsafe page. This detection (occurring for both preventative and Reactive methods) feeds through into Verification reporting.

The remainder of this section will look in-depth at these three classes of verification before looking at how the definition of verification is adapting and changing the way Advertisers look at their reporting.

Publisher and Placement Quality Indicators

Reactive verification works across a series of quality indicators. When these indicators mount up against an individual placement or Publisher they might paint that placement in a negative light (in which case the trafficker should consider having the Publisher target away from the page, domain, geography or device in question) or in a positive light, in which case more impressions should be purchased against such a placement in future. In this way, verification begins to eradicate wastage as Advertisers and Agencies use their collected verification data to map out and benchmark the safest and most valuable placements in the Publisher ecosystem to match the needs of their brands and campaigns.

Different brands and campaigns will have different creative messages and different target audiences. This means that what might be considered a negative quality indicator for one brand may not be true for others. For instance, an ad delivered to a news site featured against a disaster news story featuring a plane crash may be dangerous for an airline Advertiser but might be seen as good opportunity for a particularly risqué financial Advertiser selling life insurance. It is important for the Advertiser or Agency to make this consideration when analysing verification reporting.

These quality indicators are:

1. **Contextual:** which includes Content Classification, Audience verification and Brand Safety

2. **Discoverability:** which includes Geographic and Device reporting, Viewability and Fraudulent tag Implementation
3. **Page Level:** which includes Domain, subdomain and URL reporting and Whitelists and Blacklists

Contextual Verification Quality Indicators

Content Classification

Verification tools which specifically analyse the content of the page are also referred to as "CV tools". CV tools operate most effectively in real time. This means that when the third party ad is served to a page; the domain name, sub domain and URL are passed into the CV tool, where-upon the tool either looks the page up in its database of ready-scraped pages or it scrapes the page in real time. The scraping process captures information on the page including words, phrases, word combinations, multimedia; including video, and digests the relationship between these elements.

These elements are then categorised so that for instance, a page-level reference to a particular celebrity mean that the URL of the page is categorised in a celebrity category. Superior CV tools can analyse the semantics of the page in multiple languages and even calculate the cultural significance of certain phrases. The outcome is a statistical analysis which places an affinity of the page towards a particular categorisation. This means that if an Advertiser paid for ads to appear next to sport related content and the CV tool reported it was shown against music content, the Advertiser could take the third party report back to the Publisher in order to have future ads targeted away from the incorrect pages and onto more relevant, promised content. Publishers have access to the same level of reporting to assist them in managing their own inventory to

avoid selling inventory that is not relevant to an Advertiser from the outset of the campaign.

CV tools are interesting pieces of kit because they do not utilise user cookie information in order to provide information about page classifications. This means that they do not fall subject to any limitations surrounding privacy constraints around cookie based targeting (see Chapter 7).

The trafficker would check verification reporting around content classification and compare it to the goals of the planner in agreement with the Publisher to appear on pages geared towards certain types of content. If there is a mismatch, the trafficker would alert the planner to speak with the Publisher about the targeting and delivery concerns.

Content tends to be classified into a finite number of categorisations which are defined by the CV tool. Until these are universally accepted categorisations by all parties around the table (Advertisers, Agencies and Publishers) and written in stone by regulators like the *IAB*, there will be disputes about the methodology because content classification.

To make negotiations more challenging, custom content categories are available for Advertisers to create and define so that reporting lays out where an Advertiser's ads have appeared next to content geared towards a certain sentiment or alongside praise about a competitor. In such situations Publishers may not have the available technology in their Publisher ad server to avoid delivering ads towards or away from such pages in the future.

Brand Safety

CV tools are predominantly marketed as protection against unsafe content rather than the interference of a content categorisation. If the ad serves to a page with content deemed to be categorised as adult in nature, display violence, have content about death or disaster news stories, the CV tool can clarify this in reporting and separate it out from standard content classification.




Figure 5.2.2 The spider from the CV tool crawls sites across the internet and harvests page level keyword information and relates the information back to the page URL. The data is stored in the CV tool database. When a placement enabled for verification reporting loads it captures the page URL and passes this back to the ad server in the ad request. The ad server then matches this URL to the CV tool keyword

database and presents the contextual information in Ad server verification reporting.

Such occurrences are serious problems for Advertisers since they cause unwanted media attention or bad press when they do happen. An Advertiser may not want the carefully positioned brand to appear alongside unsuitable content. Adhering to such brand guidelines, the CV tool allows the planner to gather the data about occurrences to inform a conversation with the Publisher. The planner would initially ask a Publisher to target ads away from such page when procuring the inventory. The CV tool will show if the Publisher has honoured this agreement.

Where occurrences of failed brand safety do occur, CV tools report the occurrence alongside the location of the problem in ad server reporting.

Demographic or Audience Verification

Publishers have long been plagued by the ability to label and sell their audiences to attract Advertising spend. The methods used by the Publishers today are not standardised by the industry but the more Publishers and Advertisers that use a system of audience measurement, the more accurate and verified the audience can be. There are two mainstream methods for conducting audience verification today:

The panel-based Audience Verification Method

ComScore and *Neilsen* with their large audience panel data can give an idea (using sampling) about what audiences are connected to which pages. As a reminder from Chapter 2 - Agency Planning; *ComScore* and *Neilsen* are providing top lists of high traffic sites but can also break these sites down into matching demographics based on their panel, a product *Neilsen* refers to as *Neilsen OCR*.

A Publisher may sell empty ad spaces on a page to an Advertiser with the claim that it matches the Advertiser's audience targeting requirements (say Males aged 18-25). To verify that the audience is indeed Males aged 18-25, users on the *ComScore* panel verified as belonging to this audience type, would need to be shown ads from the tracked ad campaign.

A *ComScore* pixel, embedded in the creative or as part of the placement code loads when the ad loads and can cross reference with the *ComScore* database to see if the user belongs to the panel, and if so, what demographic they fall into. Here the trafficker needs to ensure they piggyback a *ComScore* pixel on to the served placement code within the ad server.

For the trafficker to obtain such a pixel and to get access to *ComScore* reporting (which confirmed that the ad was delivered to the right audience on the panel) would involve the Advertiser having a separate relationship with *ComScore* and paying extra for their services.




Figure 5.2.3 When a user arrives on Publisher page, it is possible to verify the demographic of the user by piggybacking a *ComScore* pixel onto a served ad. When

the ad loads, the ad detects for an ad server cookie in the usual way but in addition the ComScore pixel looks for a ComScore cookie. Relaying this information back to the comScore servers, the Advertiser can pull a report from ComScore to determine what demographic the ad was served to, providing the user made up part of the ComScore panel. Despite the sampling of the panel, across the scale of the whole ad campaign it is possible to determine an average demographic across the campaign and across each Publisher.

The social profile-based Audience Verification Method

Another methodology utilised by the verification provider *Adagoo* provides a demographic verification where page content can infer the demographic of the audience to that page. *Adagoo*'s spiders crawl public facing social profiles and mine age, gender and other audience data. In addition to these data points the spiders look at links posted on the walls of these profiles, crawl to those pages and mine the keywords and content available there. The content of those pages is then mapped to these demographic groups.

Such a system can then process the content of any page by being passed the URL of that page and report on the statistical likelihood of that page being viewed by a certain demographic group based on the already gathered data. The catch is that this is a cookie-free technology and it is not actually verifying the actual users that arriving to those pages but is rather reporting on the demographics of users that should be arriving at those pages according to the methodology.




Figure 5.2.4 The Adagoo method to deduce demographic is a little more complex. The Adagoo spider (1) crawls to a public social network profile page and extracts the Age and Gender information there (2). The spider then crawls to profile to a posted URL and goes to that URL. The spider then gathers keyword information from the page (4). The Age and Gender (2) are inferred to the keywords and stored in a demographic keyword database (5). This happens across millions of profiles and pages to build up the database. When a user (6) visits any Publisher page, the ad served placement on the page captures the semantic information on the page before serving an ad (7). The ad server passes the semantic keywords it finds, related to the ad, placement and site to the demographic keyword database (8). The database returns into ad server reporting, the inferred Age and Gender to match the keyword.

In both situations, additional code (either from Adagoo or ComScore) needs to be added to the placements by the trafficker when they are setup. But there is no solution available today to wholly verify audience information at scale. The only methods currently available entail a complex process of trying to map delivery information with extra audience data (available from data providers) and match it at the cookie level inside a DMP (see Chapter 8).

Advertisers should liaise with Publishers before the campaign begins to set expectations about the usage of Audience verification tools. The classification of content and pages must be agreed and accepted by both parties before any negotiations are made regarding those ads that were effectively wasted. If both parties are in agreement in terms of content and audience classification, Advertisers may either be reimbursed by the Publisher for ad wastage or the Publisher can use the same reports in the ad server in order to avoid wastage altogether by identifying it early on and blacklisting the pages.

Discoverability Verification Quality Indicators

Geo Reporting

Probably the most well established verification technology is Geographical Reporting (also called "Geo-Performance"). Third party ad servers can detect the IP address of the cookied user. The IP address can then be run through a mapping database which maps IP addresses to their location, the same method used for geo-targeting.

Third party ad servers will keep an up to date list of IP mappings and so are able to provide reports that show where the ads were delivered to at the country and in some instances city, DMA and zip/postal code levels. Just as with

geo targeting, region accuracy declines sharply the closer that the data tries to get to pinpoint user exact location.

Geo performance reports show the exact numbers of impressions and clicks delivered to these locations. Those campaigns intended for display purely to an audience based in the UK for instance and are shown in Geo reporting to be distribute worldwide create means for the Advertiser to approach the Publisher or network to demand the campaign be adjusted to match the right geographic audience.

Geo performance can also be used in conjunction with conversion data, to identify which regions or countries are responsible for the majority of conversion events across the campaign. Those regions found to be more lucrative could then be geo targeted with a specific creative message.

Since the Publisher can control the IP delivery inside the Publisher ad server to correct the problem of incorrect geographical delivery, third party ad server reporting is directly actionable. For the trafficker it is usually just a case of enabling the collection of IP data at the Advertiser level within the account in the third party ad server interface before the campaign goes live or during delivery. To view reporting, the trafficker needs to only look into the reporting interface in the third party ad server.

Device Verification

Browsers make publicly available their version number (referred to as a "User Agent") which can be captured by a third party ad server via a JavaScript call. The version number reveals some indication about the device type. For instance the safari browser is only supported on IOS (Apple) devices and products, therefore the presence of Safari for any version verifies that the user is using an IOS device.

Owing to the vast number of devices in use and the flexibility to install different browser versions to different devices, this verification method is more of an indicator of device type rather than make and model of the device in question. The third party ad server needs to have a general mapping database to understand which version maps to which device in order to display the results in reporting.

Device collection is becoming more standardised as part of third party ad server verification reporting, the trafficker can simply look at verification reporting inside the ad server interface to see what device types ads were delivered to. If it is found that HTML5 ads for example, are being delivered to unsupported devices or non-targeted users, the Publisher may not have installed the ad onto an HTML5 specific entity, domain or application.

Viewable impressions

The most commonly discussed verification metric among Agencies and Advertisers is Viewability. Viewability describes how much of the ad is in view for the end user and viewability is measured on viewable impressions. Viewable impressions at a minimum are impressions served above "the fold" of page as it is viewed in the users the browser. The design of most web pages contains content which fills the active browser window and trails off below the bottom fold: hidden from view. Scrolling down the browser using the scroll bar, reveals the hidden content and content disappears at the top of the browser. Strictly speaking each of the sides of the browser window are 'folds' as the content can disappear behind any of them if the page is designed to scroll in those directions.

The viewable impression count is triggered by a JavaScript embedded in the ad code which checks how much of the ad

is in view and for how long (the same technology can be used to reveal if the ad is served into the active window or tab). Some Publishers offer their inventory classified by whether the placement appears above or below the fold, and price the inventory accordingly. Increasingly in the digital marketing space, viewable impression complexity is causing advancement the capabilities of the verification measurement technologies.




Figure 5.2.5 Viewability as a form of verification. When the user loads the Publisher site some ads will be served above the fold of the user's browser while some may load beneath. The ad server delivers the ads and counts standard impressions. In the ad server interface the trafficker defines the threshold for a viewable impression (duration of exposure of a percentage surface area of the ad). Under the definition shown, Placement A would be reported as a viewable impression, as would Placement B, Placement C would not qualify as a viewable impression. If the user scrolls further down the page before closing it, Placement C may also count a viewable impression.

As the IAB develop a fixed standard for viewable impressions; Advertisers are seeking the ability to have a flexible definition of "viewable impressions". These definitions consist of a percentage of the surface of the ad exposed to the user, for a set duration of time. User

impressions that match a definition such as 90% of an ad seen for 5 seconds and above, fall into this categorisation of viewable impressions. This is called a "viewability threshold".

Ad server reporting reveals the total number of viewable impressions against each placement and when combined with media cost data, can show the Advertiser how much media spend they are wasting on ads not in view. Some Publishers will have contractual agreements with major Advertisers and big Agencies where the media cost for a viewable impression inventory reflects a standardised definition between the two parties or the definition set by the IAB. The third party ad server enables a customised threshold to be set by the trafficker at the Advertiser or Campaign levels and the data is collected and presented based on the definition.

Viewable impressions are under high demand from an industry that believes that this new metric will challenge the usual definition of standard impressions in the future. Some believe that since ads that are not being seen (according to the viewable definition) they should be dis-associated from post impression conversions and subsequent attribution. This presents interesting questions for rich media content where audio is automatically enabled; ads that are heard and not seen are still a grey area for digital measurement. In addition to the ad servers providing a record of viewable impressions, point solutions are also available for measuring viewable impressions such as *Alenty*; who bring to the table a more "dynamic threshold" methodology (where the threshold does not need to be set before the campaign begins, but the trafficker can choose the definition when the report is pulled).

Viewability is harder to measure than content categorisation because it requires JavaScript which means it is more complex to execute on Instream and Mobile ads. Viewability is also much harder for the Publisher to control and many Publishers are well aware that a fair amount of their inventory may never be seen but won't be aware of it until the ad is served. Such Publishers might argue that their pricing for the inventory is adjusted to reflect the possibility of wastage from viewability which is a conundrum for the Advertiser unless viewability can be guaranteed in the IO. To assist in eradicating this wastage some technologies are being developed that only make a request to the ad server once the placement is coming into view, rather than loading the ad tag when the whole page loads.

Double Implementations and Fraudulent Serving techniques

The publisher is not always aware that wastage may be occurring on their site, false implementations that result in poorly targeted ads (to the wrong geography for instance) do occur, especially when the operations team are working with thousands of ad tags every day. Pressing the wrong key can mean implementing the same ad tag twice in the same place and even removing a tiny piece of the ad tag code will affect the delivery of the ad which can result in a double implementation or no tracking at all (all common reasons for Publisher discrepancy).

Although these mistakes are accidental some Publishers (particularly in the long tail exchange space) fraudulently adapt ad tag implementation on purpose to serve the ad in their favour. For example if a Publisher only had a handful of pages they could double their available ad inventory by double serving ads. This way the inventory appears to be served at a rate which matches the Publishers promise (that there are twice as many pages to serve to). Lies and deceit

never work in the Publishers favour and increasingly techniques such as these are captured by CV tools and third party ad servers resulting in Advertisers stopping spend with those Publishers.

Other techniques include layering ads on top of each other, hiding ads in spots on the page that are not seen (but where code has been used to manipulate the detection of a viewable impression) and loading ads into layers which are hidden to the user. Advertisers should pay particular attention to lucrative ad spots such as for Instream ads where the opportunity to manipulate the data for the purposes of generating higher earnings is a motivation for greed.




Figure 5.2.6 *Fraudulent serving techniques see the Publisher implementing the ad in such a way so as to perform less in*

favour of the Advertiser and more in favour of the Publisher. Some ads are served into tiny iframes so that the ad cannot be seen but still registers an impression (1). Some ads are cleverly layered behind other ads, ensuring the impression is counted as a viewable impression even though the ad cannot be seen (2). The same placement is loaded up onto the same site at the same time as a copy of itself so as to serve more impressions on fewer pages (3). Instream video players set to auto-play, running below the fold, the video metrics all fire off but the ad is never seen so the Publisher still gets paid once the video has completed playing (4).

Page Level Verification Quality Indicators

Domain URL Reporting

As the ad server is able to capture the URL of the page on ad delivery, this information can be sent back into ad server reporting. By returning a full list of these URLs or even just the top ranking sub domains, the Advertiser has peace of mind that the specific ad spots that were purchased in agreement with the Publisher are the ones the ads are appearing on. Despite the fact that some buys may be across "blind networks" (networks where inventory is cheaper because the buyer buys the Audience and does not know the end page the ad appears on), ad servers can still capture and report on the page URL.

Contractually the Advertiser should ensure (via the signing of the IO when the media buyer buys the ad space) that there is no room for the Publisher to intentionally waste the delivery of ad impressions for their own gain. This is particularly prevalent in the DSP and exchange space where long-tail Publishers have a very disjointed relationship from the Advertiser and "ad brokering" is common place.

Other uses for Page Level information

Page URLs which appear in Cookie Level reporting (see Chapter 5) allow the Advertisers more insight into the CookieID itself and therefore the user. This is because a browsing history allows an inference to be made about interest categories for that very specific user (no two browsing habits are entirely the same). Aside from being useful for verification this can be utilised for targeting and even establishes itself as a unique pathway which can be transformed into part of a unique identifier for a user. This methodology is being tested for cookie-free technology in the privacy space (see Chapter 7).

White Lists and Black Lists

Across the internet Whitelists and Blacklists are enforced to certain rules. Websites can end up on Google's blacklist because of SEO violations and the rules which determine Google's indexing of websites consult the blacklist before displaying results. Verification technology can be viewed in the same way: those keywords and terms on a page that violate brand safety are considered blacklist terms and those keywords that are well matched and even contextually improve the performance of the ad on the page can be considered white list terms. Advertisers can keep record of those pages, subdomains and domains which violate brand safety and issue these to the Publisher in the IO process as a blacklist. A Publisher network may also categorise certain pages on their own networks as belonging to a blacklist, such that these pages are never sold to Advertisers for their ad space.

The URL, domain or subdomain is essentially the marker by which verification can be established. At scale, smart Advertisers and Agencies will begin to use this data to build

up a picture of the quality of the whole of the premium inventory available on the web and use this data to benchmark sites and placements against each other. This will help determine buying and planning strategies for future campaigns.




Figure 5.2.7 Placements checking for semantic and content verification can pull in the domain and URL of the page the ad is being served on. This information is fed through to Ad server reporting where Advertisers can cross reference the delivery against any existing blacklist they may have.

The combination of these quality indicators summarises the most popular methods of verification, allowing the Advertiser to check they are getting what they paid for. But in some instances the Advertiser wants to go a step further than reporting to actually screenshot the ad failing the verification quality indicators. Although not a standardised feature, some third party ad servers have crawlers technologies which can screenshot these pages and save them for the trafficker to view in the third party ad server interface. Such technologies do not always catch the ad in the act (such as with geography or device) so in the future we may see

adaptions that move from a screenshot to a short movie clip of the ad served into the Publisher page failing Verification.

See through, View through and Recordability Rates

Despite the availability of such a range of indicators, many of these will fail to report if the ad tag is loaded up in an environment which does not support the capturing of this data. Some Publishers (particularly in the exchange space - See Chapter 8) will serve ads from within an HTML iframe. Although verification works from within an iframe in most instances, iframes can be nested within other iframes. The effect is that the user sees the ad still but the tag is actually sitting in an invisible set of frames. Bursting out of multiple frames is crucial to reading the viewability of the ad and the context it is placed into. A failing to read viewability is a failure in the "recordability rate" of the ad, while a failure to read contextual verification is a failure in the "see through rate" or "view through rate" of the ad. In ad server reporting these numbers are presented against the verification data. If the rates are less than one hundred per cent then only a sample of the impressions have verification data against them. Although this transforms reporting from a full-picture into a sampling, it is still possible to determine the verification quality of placements and sites.

The Publisher has the control in most instances to prevent the nesting of iframes and in the interest of Advertiser transparency can do their best to remove such nesting in response to an Advertiser or Agency query based on the ad server reporting. Publishers which retain iframe nesting from the Advertiser perspective look like they have something to hide without such transparency in place.




Figure 5.2.8 Deep nested iframes affects both the recordability rate and the see-through rate for ads. If an ad is served onto the Publisher page within too many iframes (produced as a result of overly guarded security or excess technology on the page) the code within the placement cannot look into the rest of the page. This affects verification quality indicators such as content verification, domain reporting and viewability.

So far this section has put a lot of focus on Reactive verification which is detecting ad verification after the ad has been served. The data generated in ad server reporting is then used to help the Publisher to eradicate the wastage on future ad delivery. However, this method does not stop the ad from appearing altogether. This is where ad blocking plays a role, preventing the ad from being served based on the quality indicators.

Preventative - Ad blocking

In Ad blocking when the placement loads up on the page, if the content is found to be unsafe or if any of the other quality indicators are breached, then a Public Service ad is displayed instead, using up the impression. Ad impressions which are ad blocked by the ad server enter an ad blocking report which the Advertiser or Agency can view and download in ad server reporting to share with the Publisher. The Publisher can then either reimburse the Advertiser for the wasted ad or commit to avoid delivering the Advertiser's ads to such pages or domains in the future.

As with reactive verification, in preventative verification, the domains, subdomains and even URLs can be extracted from most third party ad server reporting and given or fed to the Publisher. The Publisher ad ops team can then add this list to a campaign or Advertiser level blacklist management facility in the Publisher ad server to prevent the delivery of the Advertiser's third party ad tags to such locations moving forwards. This process is not usually undertaken by the trafficker on the Agency or Advertiser side but may sit with an additional service team setup to engage the Publishers in a remuneration discussion: pulling the data and having such conversations. Unfortunately such negotiations can take time and the campaign in question may end before a Publisher's new billing cycle which means that recovered media costs cannot be injected back into the same campaign.

The most effective method for eradicating ad wastage in digital advertising is to use Pre-emptive verification.

Pre-emptive - Verification in Programmatic

At the time this book is being published, Pre-emptive verification technologies only work in the DSP space (see Chapter 8) and involves *not* buying or bidding on inventory because the page the impression does not meet the quality indicators for verification

defined by the Advertiser. Pre-emptive verification collects page level information about the placement, site and domain, informing the bidding process before the ad is ever served. This is much more difficult for viewability but works well for almost all other quality indicators.

Before the development of bidding technologies, the networks and Publishers were held on the promise to deliver ads into a safe environment and commit to a legal agreement with the Advertiser through contracts like the IAB supported *iASH* initiative. As a formal agreement, those Publishers that failed to comply and ended up delivering the ads into an environment that fell outside the agreement were told could lead to fines and physical auditing. In the UK, *iASH* has been disbanded in favour of new initiatives currently being developed by the IAB and supporting bodies, of which CV tools are a recommendation but as yet are not a formalised standard.

The Verification landscape

There are more point-solutions available to the Advertiser for reactive verification than both preventative and pre-emptive technologies. Reactive verification point solutions include *Double-Verify* and the technology formerly called *Adsafe* (now named *Integral*) such tools include an ad blocking feature as well. Just as with survey and research technologies, verification point solutions required an extended campaign setup as an additional line of code needs to be added to all placements in the third party ad server pointing at the technology in the point solution. Increasingly research panel data suppliers such as *Neilsen* and *ComScore* are investing in verification technologies to improve the validity of the demographic and audience data they supply. For pre-emptive verification tools are available separate from the ad server which plug directly into the popular DSPs such as *Peer39*.

The line between verification, delivery and performance can be difficult to distinguish, particularly with

around the definition of impressions and certain interaction metrics like video completion rates for Instream (which some ad servers label up as verification metrics).

Delivery and Verification reports are useful in ensuring that things kicked off okay. Once the data begins to gather in volume, the opportunity arises to conduct a more thorough analysis of that data. Beyond the most basic standard reporting templates, there are a series of more advanced aggregated reports that allow the most insightful and cost-saving trends to surface. In the next section we examine a series of these advanced aggregated reports, ensuring that live campaigns are meeting the success criteria defined by the marketers.

Chapter 5 - Section 3

Advanced Aggregated Reporting

- Cross Channel analysis
 - Day Part Analysis
 - Cross Channel Reports
 - *Figure 5.3.1*
- Advanced Conversion Analysis
 - Time to Convert
 - *Figure 5.3.2*
 - Engagement Mapping Factors report
- Reach and Frequency Reporting
 - Unique Reach
 - Average Frequency
 - *Figure 5.1.3*
 - Effective Frequency
 - *Figure 5.1.4*
 - Reach and Frequency Report
 - *Figure 5.1.5*
 - Adjusted Unique
 - *Figure 5.1.6*
- Advanced Site and Creative Analysis
 - Unique Site Overlap
 - *Figure 5.1.7*
 - Ad format reporting
 - Global Benchmarks
 - Advertiser Specific Benchmarks
 - *Figure 5.1.8*

Ad server reporting is a gold mine of data. Based on the foundation metrics of impressions and conversions, comes a whole series of answers about how to calculate effectiveness and pinpoint which impression and engagement are offering the most value in an Advertiser's digital campaign.

The following section looks at the more advanced reports available from a third party ad server reporting interface and explains their uses to the Advertiser, beginning with reports related to cross-channel analysis, before moving onto advanced conversion-focussed reporting, followed by reach and frequency reports before concluding with site and creative-focussed reporting.

Cross Channel Analysis

Day Part Analysis

Distribution of impressions, clicks and conversions can be viewed on a timeline and used to expose trends in user activity for segmented audiences at different times. These reports will expose spikes in activity that may fall in line with internet usage expectations (such as increased performance for the commuting city worker demographic, in both the lunch hour and after the evening rush hour. Day part analysis can therefore be used by the trafficker to update the time based rotation of creative messaging for that particular audience (when combined with a form of targeting such as Publisher-keyword targeting to expose the demographic of the user). In addition, day part analysis can be used to see if impression delivery at a particular time is more effective at driving conversions across certain audiences than others. This will inform the way the planner and buyer purchases ad space, since some ad spots may be more valuable to accessing their target audience at different times. Since impression volumes will be shown in day-part

analysis to fluctuate throughout the day, it might be an idea to use this evidence to justify paying more for a time based delivery, cemented into the IO from the Publisher as a commitment to that delivery.

Cross Channel Reports

Third party ad serving collects data across each channel for the purposes of de-duplicating the users (unique users seeing ads across many channels and Publishers will produce a complete pathway to conversion). In addition to this benefit is the benefit of being able to compare the performance of channels as a whole. Cross Channel reporting allows the trafficker to identify which channel is adding more value compared to others and if one channel is causing an improvement in the performance of other channels.

Cross channel reporting allows the Advertiser to recognise the value of a channel that might not directly be causing conversions, but the presence of which in a user's pathway to conversion may statistically be causing uplift on the last-click conversions of a more dominant channel.

The more commonly used "search display synergy report" demonstrates that although paid search appears to drive the most conversions under the last ad model; the display channel when paired with the search channel, creates more of an uplift in conversions than the search channel does on its own. As more channels get added to the digital marketing mix and get de-duplicated by the third party ad server; more of this style of reporting will be in demand, to ensure that there aren't channels inadvertently acting as a counter weight to the performance of their sibling channels.




Figure 5.1.1 An example of how a channel by channel uplift report adds value. If a channel is considered on its own, by just looking at a basic aggregated conversions-by-channel report, it will appear to be driving a limited number of conversions on its own. Here, where the Search channel is the last touch point by users before conversion, the Search channel is just as effective without a user encountering a Display channel ad beforehand. There are still 20 conversions, so display is not having uplift on the Search channel. However look at users converting from the Display channel as the last touch point and the analysis shows that if users encounter search ads before seeing a Display ad, there is a significant uplift. Here, the Search channel is causing a three hundred per cent uplift on the Display channel.

Channel by channel uplifting reporting allows the trafficker to clearly see how much assistance one channel is providing to another and the planners and buyers can use this data to plan when to run channel-specific activity. The same report can be used to get more granular, calculating precisely which Publishers and placements within these supporting channels are deemed responsible for the up-weight.

Advanced Conversion Analysis

Time to Convert

Setting the correct conversion window for a conversion tag is not an easy task since it is prone to fluctuations in terms of what the optimum conversion window settings should be in order to get the maximum number of conversions. A Time to convert report is an advanced aggregated report that plots the conversions on a graph according to temporal increments such as days or hours. At a point, for a particular conversion tag and corresponding campaign, the conversions drop to a level where they trail off in a long tail fashion. Chasing this long tail by cranking up the conversion window may mean leaving that window open too long and allowing ads in a campaign to become responsible for conversions where the conversion window length exceeded a reasonable consideration period. In particular this style of report can be useful to show where additional conversions would have been reported, had windows been increased. Frequently pulling this report keeps the conversion window settings in check.

Time to convert reporting, should also consider advanced attribution modelling. In addition to showing the average time taken for a user to convert via impression or click based conversion, the report can be adjusted to show how long such a conversion would take under different models. For example, with a custom attribution model that awards a greater share of conversion to clicks on rich media ads, the time to convert report would show the amount of time it takes for a user to click on a rich media ad in the conversion pathway, to finally converting. Although there may be other clicks later in the pathway, the report would only take into account, the click-to-rich-media attribution model, to advise

the optimum conversion window to exploit the maximum increment of conversions for this type of attribution.




Figure 5.1.2 A time to convert report shows how long it takes users to convert after being exposed to or click on an ad. The diagram shows that at point (1) the tipping point for the total number of conversions occurs at the 4 day mark. The trafficker could choose to set the conversion window at this point or extend it out further to catch more conversions. This chart also maps on a consideration for e-conversions, showing how long it takes users to convert under an alternative attribution model. At point (2) is the conversion tipping point. To satisfy both models on the same conversion tag, the trafficker might choose to set the conversion tag to sit between these two markers.

This advanced aggregated report includes data from all touch points that occurred within the click or view window for each converting user so as to make a consideration for

the optimum conversion window for e-conversions alongside standard conversions. The report will allow the trafficker to ensure that the windows are kept wide or narrow enough to correctly report on conversions under this alternative attribution model.

Engagement Mapping Factors report

Basic aggregated reporting can be used to show the total number of conversions and e-conversions that each placement and site might increment under the selected attribution model. *Atlas* specifically offers an advanced aggregated report which looks back at old data and applies a new attribution model to it, to calculate what the e-conversions of placements and sites would have been, had they been run under that new model. In addition to this, new attribution models can be invented and tested against existing data to see which model, provides the best results. Although this kind of reporting may seem to be manipulating the output conversion data, the models created might be a better real-world representation of how conversions are accurately attributed.

An engagement mapping factors report is useful in deciding what kind of attribution model to apply to future campaigns, providing a trial and error report output from the third party ad serving interface without affecting the originally processed core conversion data.

Reach and Frequency Reporting

Unique Reach

The concept of 'uniqueness' is an important consideration in ad server reporting. Impressions reported in the basic aggregated reporting "impressions" column are really gross impressions (to steal a term from the financial world) where as some Advertisers are interested in actual impressions or

to be more specific: Unique Impressions. The number of unique impressions will be significantly lower than the number of total impressions, since it will only count the number of unique users or cookieID's that are exposed to an ad. In third party ad serving this is called "Reach"; the total number of unique impressions. Unique reach can be controlled by the media buying process (requesting from the Publisher to only purchase space based on its unique viewing) and the creative targeting process. Creative targeting is used to ensure that either the previously encountered users are not shown any ad by the Advertiser, or are shown a different creative from the creative that they have already seen. This can easily be setup by the trafficker using ad server retargeting (see Chapter 7).

Unique reach reporting can also be setup to see if the campaign is being delivered to a growing new audience over a set period of time. The report can extend beyond the insight of new unique users being delivered ads to also see if existing users are becoming new clickers, engagers or converters as a result of adjustments to the media plan to obtain new reach.

Average Frequency

Another of the most common metrics is the "frequency" or the number of impressions that saw an ad more than once, more than twice and so on. At an ad level inside reporting the trafficker can include the column "average frequency" for an ad or creative. This will show how often unique cookieIDs are being bombarded with the same ad or creative. This is an important metric since a high frequency can have a negative impact in campaign performance and truly tarnish a brand image, while a low frequency can cause the campaign to have a lower impact. The frequency can be controlled by the frequency cap setting at the ad and audience levels (as well as on the Publisher ad server by the

Publisher ad ops team) and keeping an eye on this metric avoids creative fatigue for the user.




Figure 5.1.3 In this example an ad is served by the ad server across multiple Publishers for a total of fourteen impressions. By utilise the metrics of Unique reach and average frequency at the ad level within ad server reporting, the same ad is seen by three separate users (identified in the ad server by having separate cookieIDs). Each user is shown the same ad a different number of times but across all unique users, the ad is shown at an average frequency of four times. A simple aggregated delivery report what not show this insight, without it, the marketer could infer that the ad was shown to fourteen different people.

Effective Frequency

Effective or optimal frequency checks at which frequency level the frequency cap should be set to optimize the number of conversions, engagements or clicks for a specific ad, placement, site or campaign (and similar to the time to

convert report). Just as with the time to convert report, the Optimal Frequency report will present an optimum point followed by a drop off. However it is not advised to set the frequency cap at the point of drop-off but instead to consider setting the cap a little earlier or later depending on the goal. This consideration should be made at both a creative and site level to determine which sites or creative might require a lower frequency to maintain the target goal. The following goals can be mapped onto optimal frequency reports so that it is clear what the optimum frequency level should be:

- Maximum number of total conversions
- Highest ROI output or CPA - determined by the extended data revenue output or the cost input data on the package level (bearing in mind that an Advertiser would achieve the lowest CPA on the first impression/frequency because any subsequent exposure per user increases the cost per acquisition).
- Maximum CTR
- Maximum engagement rate or dwell
- Maximum number of specific interactions
- Optimum uplift in conversions in a separate channel (by combining with channel by channel uplift reporting)
- Optimum unique reach




Figure 5.1.4 What is the most effective frequency to hit both maximum conversions and the highest click-through rate? At a frequency of five (1) a high number of conversions are attained, after this there is a drop off. But rather than set the optimum frequency at five, the click-through rate drops off before that with the average optimum click through rate for the total number of impressions at around a frequency of three (2). To reach a balance between the two goals, the trafficker may use this report to suggest the Publishers set the frequency cap at four.

This last goal consideration is the most common pairing for optimal frequency. When the two metrics of reach and frequency are combined, an advanced aggregated report emerges: the Reach and Frequency report.

Reach and Frequency Report

The two metrics together are inversely related controls for ad spend because with a limited advertising budget only a limited number of impressions can be shown. This means

either showing the same ad to total users just once or showing the same ad many times to a smaller group. This report is aggregated to show the number of new users reached at a given level of frequency.

Reach and Frequency reports need a larger data sample than a single day in order to perform the right calculations. Typically reach and frequency reports cannot be pulled daily from the ad server but rather can be pulled weekly. With a full week's worth of data across the campaign, better optimization decisions can be made by the trafficker since it will typically take several days to establish a baseline from which to diminish frequency or improve reach. If new goals are set to increase unique reach, the report can be used to see if at the same or lower frequencies, the campaign goals are being affected.




Figure 5.1.5 Since Reach and Frequency have an inverse relationship the first frequency offers the highest unique reach and additional frequency hits the initial diminishing

unique audience. If the frequency cap was set at two (1), the impressions spent on hitting the same audience at greater frequency, could be used to establish greater unique reach on the first and second frequency.

Adjusted Unique

In Chapter 1, the problem of cookie uniqueness and cookie deletion was explored. Since cookie deletion cannot be avoided, the IAB has approved the use of a metric called "Adjusted Unique" which counts the number of unique impressions to calculate unique reach while accounting for a percentage of users that delete their cookies. In particular, the DG/Mediamind ad server allows analysts to pull this metric where-ever unique reach is used, such as in combination with frequency or determine a more realistic view on total unique conversions.




Figure 5.1.6 Adjusted unique, is a metric that takes into account cookie deletion. Here, ad serving reporting analyses all cookie to find those cookieIDs that have a longer shelf life, spanning several campaigns. This activity is then taken as a sample and the results scaled up, based on the limited pathway details available for such cookies as Cookie 2 shown in this diagram. The data from Cookies 1 and 3 is used as a

yardstick within the Adjusted Unique algorithm to predict results beyond cookie deletion and multiple device usage. The metric is adjusted for changes in cookie deletion activity over time and by user geography.

Advanced Site and Creative Analysis

Unique Site Overlap

The strategy of the Agency planner buyer is to optimize the display of ads to users in such a fashion so as to get the best response. Reach and frequency reporting gives the Advertiser insights into the relationship between the overall campaign reach and the frequency at which users are exposed to ads and creative. The site overlap report informs the planner buyer about which audiences overlap for particular sites. Since the aim of buying on multiple Publishers and Networks might be to extend the unique reach of the campaign to encounter new users, having sites which heavily overlap one another is useful to boost frequency but can damage reach. Site overlap shows how many users belong to more than one site or are unique to a site and is calculated by the ad server based on the movement of unique cookieIDs cross site and channel.




Figure 5.1.7 Unique Site Overlap reports take the insights attained from a cross-Publisher de-duplicated third party cookie and allow planners and buyers to determine if they can reach the same eyeballs with fewer sites. The three unique audiences are shown (1, 2 & 3) revealing a clear overlap. The report indicates that Site D can be omitted from the plan, if the plan is to reach a unique audience but utilising the impressions delivered from Site A and Site B.

An additional consideration is whether spend to reach users on a previous campaign are affecting this one. Although the Advertiser may be advertising a different product, the old messaging may still have had some impact, potentially adding to user exposure in the form of frequency. In this way, the spend on those placements may be contributing to the performance of the new campaign. Pulling a report which layers unique reach and frequency, cross campaign, channel and site, allows the planner and buyer to consider reducing spend after a high spend campaign, owing to the marginal impact it may have on the upcoming campaign.

Ad Format reporting

Although it is possible to pull basic aggregated reports for delivery on a campaign level by Publisher and to show the best performing ad format, it is more challenging to just see if one format outperforms the others in every situation. Some ad servers make this easier for the analyst in the third party ad server reporting interface by aggregating click through rates, effective frequency, effective improvements in unique reach, total impression delivery, conversion rates and effective ROI lead by ad format. Such reporting is useful to justify spend for Rich Media for example and is most useful for Creative Agencies and planners in deciding what dimensions of ads to buy and develop.

Global Benchmarks

The urge to look over a friend's shoulder in an exam is not always to cheat but rather to check that your answers look right. Third party ad server benchmark reporting takes campaigns results from all Advertisers that are opted in, to aggregate those results. Results are broken by country to make them more relevant to analysts. This mass aggregation allows Advertisers to see which sites, placements and formats are the most effective before running such activity themselves. The aggregation is fully anonymous so that it is not possible to determine the source of the data by analysing the results. This would normally be possible if an Advertiser from a particular vertical was the only Advertiser in that vertical to either be utilising a particular Publisher or be saturating digital advertising in a particular market.

To make the data more relevant, ad serving providers compartmentalise the results by Advertiser marketing vertical so that Advertisers in the Financial vertical can see how well 300x250 expandable ads may perform as a CTR on specific section of a specific Publisher site based on the benchmarking results from the previous year. Conversion rate data is also available, but since conversion tags can be used on all manner of Advertiser pages, the ad server will group together all conversion tags labelled as "sales tags" to determine aggregate conversion rates.

Typically ad serving providers will offer benchmark data at a very high level to all Advertisers and Agencies using its interface while offering deeper insights to a different set of benchmark data aggregated by specialist opt-in Advertisers and Agencies, who only get deeper insights by offering up more of their own data to the pool. These deeper levels of collaboration are tricky to operate, they require a blanket level of consensus among all participants which in practice is complicated to maintain. Some participants to the data set

may feel that they are giving away more than they are getting back and convincing contributors that there is universal fair access is challenging to prove, especially if Advertisers differ in size. To provide a fair balance, there is usually a minimum contribution requirement for reporting. This can mean that despite the huge amount of data available in ad server reporting, the benchmark report may display blank results where only one Advertiser in a country has any ad serving data for a 300x250 expandable ad on a specific Publisher site.

Advertiser Specific Benchmarks

Large Advertisers and Agencies also generate such benchmark reporting outside of the ad server, by feeding their ad serving data into a separate reporting system where they can layer in cost data and more specific performance data. The cost data would also be a record of what was paid previously to determine which sites, placements and formats that are the most cost effective for their expected performance. Such benchmarking information would likely feed directly into an Agency or Advertiser planning tool or Agency trading desk to better inform the planning and buying process.



Figure 5.1.8 Benchmarking data is universally available from ad servers, fully anonymous showing a breakdown of performance benchmarks per vertical and ad format. Specific Advertisers could also opt in to a pool where more of their data contributes to a better indication about performance data down to a site or Publisher level. Beyond this large Advertiser run enough activity across Publishers to allow their planners to view specific cost and performance data by format and placement. This informs the decision to buy those ad spots again in the future but typically this is not information that is shared with the Publisher.

This data is not just useful for future planning but also to compare with existing performance to check that the campaign is on track. Numbers reported to be an improvement on the benchmark are a positive story for the Agency analyst who can use the benchmark to set expectations with the Advertiser or marketer for the reported results.

Advanced aggregated reporting is a method to pivot and adjust the way that data is viewed, building on top of basic aggregated reports. Advanced reports are useful because they provide a way to make more sense of the data without diving deeply into the cookie level reporting, also available from the ad server. The next section looks at more niche channel-specific reporting, which is useful in silo, allowing marketers heavily focussed on Search or on Video, to just look at the delivery, engagement and performance of those channels in isolation rather than looking at the bigger picture.

Chapter 5 - Section 4

Channel Specific Reports

- Display Channel Reporting
- Rich Media Channel Reporting
- DCO Channel Reporting
 - Multi variant testing
 - Dangers in Version Overwriting
 - Version reporting
 - *Figure 5.4.1*
- Instream Channel Reporting and Metrics
 - Video Started and unique video starts
 - Ad duration
 - Increment ad play-through (25, 50, 75, 100%)
 - Average % video completion
 - VPAID custom interactions
- Mobile Channel Reporting
 - Cookie Duplication
 - Cookie Absence
 - Device Finger Printing
 - Cookies in applications
 - Device and Carrier reporting
- Paid Search Channel Reporting
 - Search Overview
 - Keyword Pathway Analysis
- Natural Search Reporting
 - Double counting and Single Counting
 - *Figure 5.4.2*
- Affiliate Channel Reporting

The third party ad server reporting interface offers traffickers and analysts reports that contain metrics relevant only to specific channels. These channel specific reports match well to the same major channels described in Chapter 3 - Campaign Management

by Channel. The following section will outline the proprietary metrics, methodologies and analytical considerations within channel-specific reporting.

Display Channel Reporting

Reports specific to the display channel (omitting rich media, DCO and video from the definition of the display channel) are the most basic kind of reporting available and would be the same basic aggregated reports described as delivery and performance reports. Here reports predominantly relate to impression information and click information for the Publisher sites, packages, placements and ads. The additional considerations, more unique to the display channel are those where the naming convention can be adapted to give more information about the creative. These fields in reporting would be "creative name", "ad name" and ad and placement classifications. Creative name and ad name, as described in Chapter 2 - should match a description of the imagery used in the creative or relate to the product advertised within the ad. Within reporting these names can be aggregated so that if they are used. The ad and placement classification fields are also useful for this purpose but serve to provide additional meta information. Ad classification one might be populated by the trafficker to contain the date that the creative was swapped out for another with the same name. Rows within reporting could then be aggregated to show the different in performance based on this information.

Also relevant to Display reporting is the ability for the ad server to reporting on backup image delivery. Sometimes a basic aggregated delivery report will hide the number of impressions where users saw a backup image instead of the flash creative, channel specific reporting should provide the analyst with the ability to see this information as well. An indication to the delivery of a large number of backup

images may mean that unsupported devices are being targeted at the Publisher level; a control which can be adjusted by having the trafficker feed this reporting information to the Publisher to target away from non-flash supported results before a high delivery of backup images equates to wastage.

Rich Media Channel Reporting

The primary difference from a reporting perspective between Standard Display and Rich Media, is Rich Media's ability to track interactions and subsequently pass them back into reporting. Interactions can be divided into two groups; default and custom interactions. Default interactions apply to the counting of the trigger of certain Rich Media specific functions such as Expansions. Total number of expansions, for instance is a standardised metric in Rich Media reports. Custom interaction tracking for Rich Media allows the creative developer to place a trigger inside the creative anywhere, so that an action within the creative will count that custom interaction occurrence. Creative developers are encouraged to work with marketers and Agency teams to put useful custom interactions within the creative, tracking anything from a user hovering over an object to clicking, pressing a button or even sitting idle.

In reporting, interactions can be tied back to unique impressions and conversion events. The overall time spent interacting within the ad (such as a mouse pointer hovering over the creative for an extended period of time) is reported in metrics such as "interaction time" or "dwell". The measure of such metrics is protected from accidental trigger by putting such safeguards in as waiting for one second to pass with the mouse pointer operating within the confines of the ad before counting begins. The key here is that Rich Media creative can also support a timer such that an analyst can pull from reporting a "total brand exposure duration"

metric which is a total count of interaction across all users. Such metrics could be used to compare with TV ad campaigns that are measured by minutes of exposure to the total audience base.

Some ad servers offer a heat mapping tool, which layers hotspots over the creative to show analysts where users were predominantly engaging or interacting with the ad. Although such reporting is aesthetically rewarding the analyst, the interpretation of such reporting can be subjective and difficult to optimize from creatively speaking.

DCO Channel and Audience Reporting

DCO reporting can be viewed at a high level; such as to view the effectiveness of versions based on the version name, but deeper insights are available by ensuring such reports display the detail of the dynamic elements of that version so that results can be tied back to the correct recipe of those elements. Multi-variant testing specifically allows for the analysis and identification of best performing versions or variants based on those best performing elements.

Multi variant testing

With the dynamic elements of the ad exposed in the creative the reporting may expose trends such as that a particular price point across lots of different types of creative always seems to pull a high number of clicks or conversions. The same could be said of a background image, colour or font type used although clearly some co-incidental trends may surface which fall to the best judgement of the campaign manager to dismiss or de-prioritise in favour of optimizing to elements which logically make attract a better performance (such as price point or offer type). In reporting for multi variant testing, the usual column for the version name is replaced with the name of the

varying asset and reporting such as impressions, clicks, conversions and interactions can be layered up by site and placement so as to find out if a combination of assets equates to a better performing version when they are combined together. Such reporting is of particular interest to the creative developer.

Dangers in Version Overwriting

Updating a price point or any other dynamic element for a specific version in the setup of a DCO campaign is a bad practice in the ad server without changing the name of the version entirely. Some ad servers do not place restrictions on version updating simply because restrictions would add extra steps for Advertisers and make the setup more challenging. However; the amendment of a dynamic element in an existing version is not a recorded event so inside DCO reporting, a top performing version which was changed at some point in the day could not attribute its success to the amendment.

Even at the cookie data level the DCO version amendment is not time stamped so the trafficker would need to make a full record of changes made to existing versions. The solution is to always create new versions when any dynamic element is updated and ensure there is some relationship between the new version and the old version to understand the version evolution (this can be done with a good version naming convention).

Version Reporting

DCO reporting within the ad server reporting interface costs of reports that shows version delivery, performance and engagement at the version level instead of at the ad level. Detailed reporting of this

type will also show which audience or placement the version was assigned to, to allow the trafficker to optimize the delivery, targeting and rotation of best performing versions. Version level reporting can become complex to analyse if only the version number was displayed for each row in the report since the number or ID does not give away an accurate description of how versions differ from one another.

For this reason, larger version reports that include a list of the content of the dynamic assets is useful although the reporting quickly begins to replicate the spread-sheet based mass version interface within ad server campaign management. To reduce report size, some ad server reports for DCO will omit rows or versions which recorded no delivery or poor performance, reducing a potential pool of thousands of versions at a by day, week or month granularity, down to just a few rows. This makes life much easier for the analyst and trafficker.

Version	Country	URL	Impressions	Clicks
901	UK	uk.msn.com	80000	170
902	DE	de.msn.com	67000	511
903	ES	es.msn.com	87000	601

Version	Background	Logo
901	red	logo2.gif
902	red	logo2.gif
903	black	logo2.gif

Figure 5.4.1 Version level reporting (shown left in this diagram) reveals the delivery and performance of the various versions. Each of the versions is seen to be running against a different defined audience type, here by geography and clicking through to a different part of the MSN.com website. Beyond this the differences between the versions are not revealed in the report. Navigating the ad server campaign setup interface reveals more information about each version by how they vary from one another. The analyst would need both these data sets to adequately analyse each version and explain whether the aspects of each version or the targeting are responsible for the performance.

Instream Channel Reporting and Metrics

The metrics that can be tracked in the Instream channel have growing synergy with TV measurement. Historically big Advertisers, advertising through the medium of television have been able to track a sample of ad delivery and exposure. This allowed Advertisers to pay for TV ads in new ways: such as only paying if the video of the ad was played in full. Apart from the Display channel and Search channel, the Instream channel is the only other channel where Advertisers are paying Publishers not on ad delivery, but on channel-specific metrics: video plays and video completions. The following is a list of notable metrics available only in Instream specific reports.

Video Started and unique video starts

At the creative level, the analyst can determine how many times the video was watched and how many unique users watched in (some users will play a video through more than once). Video started is useful to compare with the number of impressions for the same ad. Since the video started playing after the impression delivery, some Advertiser may prefer to pay Publishers based on the video started metric.

Ad duration

Ad servers can provide meta data about the video creative itself such as the quality of the video stream (in bits per minute) and the total length of the video. This information can help analysts perform side by side comparisons between different video ads which may be of differing lengths and quality.

Increment ad play-through (25, 50, 75, 100%)

The amount of time that the user views the video ad for, is another Instream metric (but one that is also recommended to track in rich media video units). These metrics will show user drop off as the video nears its completed play. A video which is too long to hold user interest would show high drop off of the total number of unique impressions as time goes by. As an action the trafficker would request a shorter video length from the creative developer and to buy shorter video ad space from the planner and buyer.

Average % video completion

At the ad level the analyst can see which percentage of users watched the video to completion. This can act as a marker for delivery, performance or verification depending on how the analyst interprets the data.

VPAID custom interactions

Just as with Rich Media, Instream video supports custom interactions. These are inputted by the creative developer before the campaign launches and can be positioned anywhere in the video creative. Buttons which allow the user to share the video on a social network or open an Instream video gallery from within the VPAID creative can have such actions tracked and incremented within Instream specific reporting.

Mobile Channel Reporting

Mobile specific channel reporting, although very similar to standard Display delivery reporting, has some considerations to make which relate the identification of a unique user. In standard Display reporting, knowing a cookieID is unique or not unique allows the user pathway to conversion to be

mapped cross Publisher for the same device. Once a second device or a mobile application is introduced into the mix, the cookie dependent technology related to establishing unique users and conversions, is lost.

Cookie Duplication

The human race now accesses the internet more frequently through a mobile device than a PC but the environments from a browsing experience across devices is unbalanced. In the PC environment, the ad server drops a cookie, labelling the user and tracking them. So long as the user only uses that same PC and browser to access the web, the Advertiser only needs to worry about cookie deletion, but Mobile presents a problem of cookie duplication.

A single user that switches between devices means that the ad server sees multiple users in its data collection at the cookie level, than the true single user. If there are multiple cookies in existence at the same time for the same user these need to be tied together to establish the true reach of the campaign and for pathway to conversion data (see the final section in this chapter). To overcome this problem a single technology needs to be used on both devices so that in reporting there is a common link between the two cookies. An additional data set can be applied to the mobile cookie level reporting to perform what is known as "Cookie Matching" to filter the report to establish the unique users.

Such datasets are available for Advertisers to buy; the marketers would then pull this data into the ad server cookie level data inside a DMP environment to cookie match. Data sets that could be used are ones where the user has logged in to the same place on several devices

(an example of which would be signing into a browser as is possible with *Google's Chrome* or logging into a cloud-based account such as an email provider). This is all well and good while there are cookies to match; but the stinger is that third party cookies are not ubiquitous across mobile devices.

Cookie Absence

Apple made a decision to have the *Safari* browser not allow third party cookies as a default setting. This has damning consequences for tracking activity on *Apple* devices. If a cookie cannot be dropped there is no way to knowing if an individual browsing user has seen an ad and if they have seen it more than once. Inside the third party ad server database, the cookie is simply given no ID. As apple grows in adoption, so the number of traceable users decreases using the cookie level method.

To try to get clever, technology providers realised that they did not need a cookie per say, but rather just needed a common identifier for the device. The identifier would need to be unique so that the user on the device could be detected as a single individual and be measured as a separate "cookie" from the others being tracked by the technology. This identifier is referred to as "Device Finger Printing".

Device Finger Printing

At the core of the *Apple* device sat the "UDID". A Unique Device Identification label which technically speaking could be accessed by browser based technology and then a matching "cookie" could be built remotely in the server environment. The effect is that the ad server would "see" a single user. But the UDID was not to be the saviour for the cookie problem that many had hoped. *Apple* scrubbed the label from all devices, well aware of the

workaround, and at time of Publication had begun presenting access to a new identifier to ad developers, specifically for the purposes of Advertising.

The purpose-built identifier has the possibility of user control so that they can opt out of cookie-based targeting by switching it off. This exposed the control *Apple* has over third party ad delivery and tracking, by blocking or controlling access to a unique identifier it is not possible to make an independent audit of the true number of unique cookies on such devices which is a problem if *Apple* decides to enter the world of Advertising Technology and supply of their own media space.

Workarounds do exist to try and circumvent the problems of identifying and counting unique mobile users such as creating a unique "session ID" on the fly, or using the IP address to establish a single cookie but these techniques remain flawed and the problem of a lack of common cross-device cookie remains in place.

Cookies in applications

Cookies are enabled through browser usage and are unique to browsers so in-app advertising presents a challenge for data collection since there is no standardised in-app cookie to speak of. The apple device identifier for advertising could be used in this instance (captured by the app using a specific line of code and then passing up to the ad server when the ad is delivered) but this still means there is a need for cookie matching.

Some SDKs force the mobile device to momentarily load up the browser on first use of the app so that the ad server can look into the browser to detect the browser cookie ID before sending the information onto the ad

server. From that point forward the device can be associated with a traditional third party cookie. This is not a technique recommended to the development community by the leading app stores but again another workaround that helps demonstrate the jerry-rigging required to data collect in a yet-to-be standardised channel.

With the knowledge that cookies remain elusive in the Mobile channel, ad server reporting for Mobile restricts the universal metrics available for the analyst to look at. Impressions, and clicks are available at a Publisher, Site, Package, Placement and Ad level but reach, frequency, cross-channel and conversion reporting are as yet unavailable or unreliable across the board. Rich Media ads within mobile will support interaction reporting for HTML5 ads and flash where it is supported. IP delivery can also be complex due to the portable nature of the device, so geographical reporting is ill advised. In addition verification has limitations for devices that do not support JavaScript meaning that at time of Publication it is very challenging to measure viewability in mobile.

Device and Carrier reporting

Some ad servers have the technology in place to allow the Publisher to return values on mobile placements which may provide more information about delivery. These include fields such as Device and Carrier information which a Publisher could pass back to ad server reporting when the ad delivers. Having said this, few Publishers will not go that step further to providing this extra information since there is very little that can be done on their part to control this form of delivery for the purposes of media optimization later on (See Chapter 6).

Paid Search Channel Reporting

Although paid search reporting is available from inside the Search Engine marketing (SEM) interface, the ad server reporting interface presents the analyst with the opportunity to see this same data alongside the data from other channels all within the same interface as well as conducting keyword path analysis.

Search Overview

Standard Search channel reporting allows the trafficker and analyst to look at the delivery and performance of keywords, ad-groups and search ads alongside the CPC cost data retrieved by the ad server from the search engines via the search engine's own reporting API. The difference is the ad server can establish conversion reporting (to ad server conversion tags) as well as look at the uplift in conversions and clicks from the search channel to other channels like display (as explained in Channel by Channel uplift reporting).

Keyword Pathway Analysis

A keyword path analysis report aggregates the most commonly searched for and clicked on search terms in their lead up to the conversion. These combinations of search terms and the time stamps between them help the analyst understand where various keywords sit in the user's pathway to conversion and allow the trafficker to subsequently up-weight or optimize the bidding on certain keywords back in the engine. This can be done by analysing the reporting or by connecting pathway to conversion feed data (see next section in this chapter) back up to the bid management tool. In addition, some keywords may not be driving conversions at all, as such the trafficker can effectively eliminate them from the bidding process for the rest of the campaign or subsequent campaigns.

Natural Search Reporting

Natural search tracking allows the trafficker to track natural search as though it was another channel in the paid-for marketing mix. However, the key to making the most of the collected data, is setting up reporting correctly. The risk is that if user pathways to conversion were analysed, most pathways would end with the click of a natural search keyword in an engine. Under a basic attribution model like "last click", natural search would appear to be driving all conversions. In standard conversion reporting it would then appear as though all other paid-for channels could be removed from plan without any affect. In fact conversions would drop off or disappear completely, so for the Natural search channel, the trafficker must ensure that full attribution is not awarded to Natural search. This can be done using an advanced method in ad server reporting to separate out the activity completely from all other channels and then rebuild the user pathway to conversion later on.

Double counting and Single Counting

To still provide visibility on the natural search impact, while not diminishing the value of other channels under the last ad model (see Chapter 4), natural search should sit completely separately from other campaign activity by having it live in another Advertiser account entirely separate in the ad server to avoid de-duplication. There is an option in the Atlas ad server, called Single Counting Vs. Double Counting, which allows for this same affect.

If the user sees a display ad, and then search on a search engine for the featured product and clicks a natural search result on the search engine results page the following happens:

- The display ad, trafficked under the Advertiser's main Advertiser account records an impression followed by an impression based conversion from the display ad to the

conversion tag on the landing page. The impression and conversion are recorded against cookieID 001 in cookie level data.

- On the landing page, the natural search script fires a click tracker, originally setup as a tracking placement under the Advertisers natural search Advertiser account. A click is recorded in this account against this separate Advertiser but also against cookieID 001.
- The same script then fires a conversion tag, setup under the natural search advertiser account, recording a click-based conversion against cookieID 001 in cookie level data.

If the analyst looked in the pathway to conversion data (see next section in this chapter) for the main Advertiser account the attribution for the conversion, under the last ad model would go to the display event, with no natural search activity recorded in sight. The analyst could then log out and log into the natural search Advertiser account. The pathway to conversion data report here would show attribution to the natural search click activity. These two reports would need to be combined at the cookie level with timestamps matched, to see the user's true pathway. What this workaround does, is that it preserves the attribution of the conversion to the display ad in the main Advertiser account without needing to create a custom attribution model or undermining the impact of the paid for activity.

Some ad servers offer SEO specific attribution considerations automatically along with a JavaScript embedded into full-site distributed conversion tags such as with a universal tagging solution (see Chapter 4) so that there is little or no effort on the part of the Advertiser but predominantly ad servers require the style of workaround

described to support natural search tracking in ad server reporting.




Figure 5.4.2 To calculate where Natural Search sits in the pathway to conversion, the setup requires a workaround which feed reporting data into two separate Advertisers. A report can then be built from the cookie level data to see the pathway to conversion while still awarding the final conversion to paid media channels in aggregated reports. This diagram shows how a single user sees an ad at 1pm. Later in the same day the user searches a keyword and navigates through to the same landing page as was designed for the paid-for ad (Ad 1). The script on the page registers conversions in separate Advertiser as the tracking is setup in the campaign management section of the ad server interface. At the cookie level, the cookie history can be used to re build the pathway to conversion across the single user.

This setup style for natural search means that within aggregated reporting, natural search behaves the same as any reporting related to click counting. This means that total clicks and click based conversions can be recorded against this channel within the specific natural search Advertiser.

Affiliate Channel Reporting

The idea of affiliate based reporting within the ad server is to determine which affiliate in the affiliate network, given the use of a single ad, is responsible for the click and conversion activity.

To calculate which affiliate site was responsible for each sale, the affiliate network generates a unique reference number for each affiliate when they generate the Advertiser campaign code before placing it on their website. When the user clicks the activity and goes on to purchase, the affiliate ID is carried all the way to the purchase page in the click pathway where the data is collected and a feed of information is returned to the affiliate network about affiliate so that they can distribute the payment to the affiliate sites. To achieve this the click through URL consists of separate parts which each report back to their various tracking technology. The combination of these parts is referred to as "Daisy Chaining".

The output from ad server reporting on affiliate activity, lists the affiliate site name or unique identifier on each row along with the total number of attributed conversions and the aggregated ROI. This is based on the extended data revenue passed through at the point of conversion on the conversion tag extended data parameter.

The vast range of reporting options means that the number of metrics possible to pull down from the third party ad server will be in excess of one hundred different values. Reporting templates which can be customized will usually allow users to bring metrics into reports which may have no place being there: such as Rich Media metrics in customized Paid Search reports. It is for this reason that traffickers and analysts using the reporting interface are urged to play around inside the interface, before pulling reports for a more urgent purpose. Without thorough practice, the act of report pulling can quickly become frustrating unless it is clear which report will report the required stats.

This section concludes the summary of aggregated reporting options. The last section in this chapter covers off the reports that are usually found outside of the ad serving interface, although the data still originates from the core ad serving database. An example of such data is cookie level data which is usually so detailed and granular that it requires much greater processing power to output and present. As such this style of reporting will usually be delivered manually to the Advertiser or Agencies FTP (File Transfer Protocol) site.

Chapter 5 - Section 5

Cookie Level data

- Reporting at the Cookie Level
 - Figure 5.5.1
- How is the data supplied?
 - Look-back Data Dump
 - Cookie Level Feed
 - Report delivery, size estimation and storage considerations
 - Metrics in Cookie Level Data
 - Tracking impulse conversions
 - Pathway to Conversion Reporting
 - Figure 5.5.2
 - Point Solutions for Advanced Cookie Level Attribution
- Reporting on Conversion tag Actions
- Extended data at the Cookie Level
 - Figure 5.5.3
 - Figure 5.5.4
- Impression and Click based token data at the Cookie Level
- What is missing from ad server reporting?

Aggregated reporting brings together various metrics to give a high level view about the total number of those metrics counted for a given ad server hierarchy entity over a given period of time.

This means that the trafficker or analyst can log into the third party ad server reporting interface and download a report by a given date range to show total impressions, clicks, conversions,

interactions by Advertiser, Brand, Publisher, Site, Section, Domain, Package, Placement, Ad, Creative and Version. To really learn exactly how users came across ads and the story behind the pathway to the conversion, it is essential to dive much deeper. In cookie level data, instead of each row of the report being a summary or activity by entity, each row is a summary of activity by cookieID or at the maximum granularity can be a summary of a single recorded action or event for a cookieID.

Reporting at the Cookie Level

Cookie level data exposes every recorded event in the third party ad server for an Advertiser as it is related to each user (or each unique cookieID) and places time stamps against various activities. In just a few rows or columns of this level of raw data, it is clear to see how and where a user saw or clicked an ad in a campaign over an extended period of time if or before they reached a conversion tag and converted. This is enormously useful for analysts and for support teams troubleshooting misbehaving ads as it demystifies the aggregated data. The simplicity of the aggregated data fails to explain how or when the aggregated data was collected and without this context, the reports can be misunderstood.

This is particularly true when asking the analyst to answer more complex questions with the data such as: What products did my most effective ads, drive users to purchase? Or is there a relationship between the final cost of the purchased product and the Search keyword? Only cookie level data can glean these answers without specialist build aggregated reports to pivot the data in ways that may be more specific to a particular Advertiser. Cookie Level data is typically available to the Advertiser or Media Agency filtered in the following ways:

- Cookies to all Impressions only (with the option of adding interactions and IP)
- Cookies to all Clicks only

- Cookies to all Impressions and Clicks only
- Cookies to all Conversions only
- Cookies to all Impressions, Clicks and Conversions (with the option of adding interactions and IP)
- Pathway to Conversion

Each filter method provides a different output to the Advertiser although all will be in a view that looks like a spread-sheet or CSV format; make it easier to conduct analysis across the data. The benefit of choosing one filter over another is the amount of data that the file will contain.

Cookie ID	Action Type	Time of Action	Advertiser ID	Ad ID	Placement ID	Site ID
00001	Impression	3.45pm 01/01/2015	8192	9675748	49495993	4994

Figure 5.5.1 an example of a single row of cookie level data filtered by impressions. Here a unique cookieID was seen for a specific ad to generate an impression, within a specific placement on a specific site at single point in time. A cookie level data report will typically contain many thousands of such rows.

Cookie level data can be vast. For each passing second, millions of users are encountering ads and then taking actions or converting which creates hundreds of millions of rows in an enormous database behind the scenes at the ad server. Cookie level data is commonly misconstrued as the raw third party ad server counting database (which is built in the moments that ads are served and metrics tracked). This is rarely the case; cookie level data is constructed from the raw data later on. The raw counting database is produced and stored separately in a much more garbled format, that cannot be read by the analyst without assistance or decryption (this dataset is often referred to as the server "log files").

Cookie level data is typically provided to the Advertiser or Agency by the third party ad serving provider at additional cost, since it

takes a lot of processing power to churn through and build the data set, before using up a lot of bandwidth to deliver it.

How is the data supplied?

Cookie level data can be supplied in two types.

1. A look-back of cookie activity over a period of time (depending on how long the third party ad server stores or has been collecting the data for)
2. A feed which supplies data from the previous day or week.

Look-back Data Dump

Third party ad serving providers do not usually collect cookie level data as a default. The collection of the data (in a format that can immediately be presented in a cookie level report) needs to be switched on by the trafficker or the ad server support team. From the point at which it is enabled is the earliest date on which cookie level data will be available as a report. A look-back data dump can provide to the Advertiser or Agency at a point in the future, a massive file containing some or all of the cookie level data. These files are usually very large and are difficult to open and read. It is recommended for analysts working with data of this size to instead opt for receiving a trickle of the data every day as a feed. This will allow them to build up a larger database of the data on the Advertiser or Agency site later on.

Cookie Level Feed

Requesting a feed of cookie level data from an ad serving provider allows Agencies and Advertiser to store the most granular data available with the freedom to run complex formulae across it at a later date (or in parallel to running aggregated reporting from the third party ad server analytics interface). Larger Agencies invest in technologies that pick up the cookie level data delivered by the ad server, process

it, and present it in customized ways. Some of these in-house technologies then automatically take actions as a result of the data such as buying media programmatically for search via a bid management solution or for Display inventory (see Chapter 8).

Report delivery, size estimation and storage considerations

Cookie Level datasets can be monstrously large in the look back form and even as a feed; the amount of data quickly begins to add up. Therefore the biggest cost is not the retrieval of the data but the long term storage of it. Despite the global cost of data storage still shrinking, these data files are too large for even the best ad servers in the world to afford to store data for all Advertisers for a lifetime. Furthermore as ad servers release more sophisticated technology and features, the amount of data that can be collected quickly and quietly grows, creating further storage demands.

Usually an ad server will supply cookie level data to a secure FTP site where the Agency or Advertiser will retrieve the data set on a rolling cadence. The ad server is therefore expected to deliver such data at the same time each day or week so that subsequent analysis from that data set can be planned for effectively. The trafficker could calculate how large the cookie level data files will be by looking at aggregated reports. The total number of impressions will indicate the total number of rows in the cookie level data report.

Metrics in Cookie Level Data

Cookie level data typically matches four key metrics of data to the ad server hierarchy entities; impressions, clicks, conversions and rich media interactions. On occasion cookie

level data is enhanced with geographical information or verification data. Analysts and marketers that plan ahead would think getting all of the available data is the wisest decision but of course this means much larger file sizes for storage. Some ad serving providers get around this issue by offering mapping files (where names in a particular column or group of columns is just repeated over and over). In such an instance a mapping file is supplied to allow the analyst to transform the repeated data, this method shrinks the file size of the larger data set and makes the data more portable. An important consideration for cookie level data is user cookie deletion. A user can be expected to delete cookies within 90 days of receiving them, so the data is not useful long term for the purposes of retargeting the same users again (see Chapter 7). Rather the data is used to surmount trends and predict behaviours for the purposes of optimization (see Chapter 6).

Tracking impulse conversions

Since basic aggregated reporting predominantly reports at a day by day level, those campaigns requiring a faster response, targeting and buying by the hour instead of by the day, require timestamps to optimize accurately. Although day part analysis can be revealing to an extent, it fails to answer questions about user actions undertaken such as "impulse conversions" which may only happen moments after user exposure to an ad. Cookie level data can reveal this information where aggregated reporting cannot.

Pathway to Conversion Reporting

The most common reason Advertisers and Agencies use cookie level data, is to understand more about the cookie-level Pathway to Conversion (P2C) also called "Exposure to Conversion" or "E2C". This granular report shows how each converting cookieID encountered any tracked or served Advertiser activity before going on to trigger a conversion

tag. The record of the total number of touch-points shown in this report for each cookieID, on its pathway to conversion before the conversion activity is typically limited, so that the file size is not too large. Smaller Advertisers are usually comfortable with a touch-point look-back of ten "events in the pathway" but the larger Advertisers, with bigger marketing teams, more budget for data analysis and complex, wider reaching campaigns may require a pathway length of more than one hundred touch-points or events.




Figure 5.5.2 Full P2C (pathway to conversion) reporting at the cookie level. Such reports are granular and very extensive. Diagram depicts just four touch points on the user pathway to conversion which in actual fact could contain many more touch points. Each touch point or event is made up of data about the Publisher page and the ad. This includes the timestamps for seeing or clicking those ads as well as ad and placement classification information. The final touch point is followed in each row by the conversion information and extended data so that for each CookieID it is clear to the analyst all the data points that were captured that in some way could be attributed to the conversion event.

Having ten events in the pathway means that any events that occurred before the last ten are removed from the dataset, since these were so far away from the actual point of conversion that they can safely be disregarded. On the other hand, if some of these events logged the user encountering particularly impactful executions, then it may make more sense to look back further or to have a more intelligent way to filter the events. This is where attribution modelling is useful, since it can filter out or play down pathway events and provide an indication as to the contributors to the conversion from within aggregated reporting (see Chapter 4).

Advertisers that run a lot of activity, saturating the online space with ads will want to look at more events in the pathway to determine what impact all their ads are having across multiple campaigns for all converting cookies. Such situations are highly complex and require analysts to write additional programs to process the cookie level data, assigning statistical value not just to attribution that a single pathway event would had to a conversion tag, but to all conversion tags across all campaigns.

Point Solutions for Advanced Cookie Level Attribution

Beyond the conversion, some Advertisers want to analyse what impact pathway events have when such events lead to non-conversion performance events such as clicks, interactions or improvements in verification and viewability. Many technologies in the space are ill equipped to deal with this level of complexity but some point solutions have emerged to address such analytical demands. These solutions include *ClearSaleing*, *Encore Metrics* and *Visual IQ*. All can be setup to re-process cookie level data out of the ad server at the maximum granularity or have their own pixels placed in ads, click pathways and

conversion tags in order to replicate the data set in these separate systems. The issue with implementing additional tracking for this purpose is that they can come at a high additional cost to serve and track the campaign for insights which may not provide the best ROI. In addition, the additional tracking technology will be counting the impressions, clicks, conversions, actions and interactions at a very slightly different time from the ad server so there will be a minor data discrepancy which must be taken into account when conducting an analysis. For these reasons it is recommended that advanced attribution point solutions take and process the existing cookie level data, returning the insights to the Advertiser or Agency directly.

Reporting on Conversion tag Actions

Granular reporting beyond aggregated reports are referred to by some ad serving providers as "Raw Data Reports" or "Data Pass Back" reports, and offer a granularity such as an aggregated by conversion tag extended data custom parameters (see Chapter 4) or a list of conversion tag "Actions"; showing the number of times that a conversion tag has fired in addition to conversions. A report of this nature can be crucial to troubleshooting the behaviour of conversion tags as it shows the number of times a conversion tag is loading and not just the total sales or counts attributed to traffic to the conversion tag, coming from an ad.

Advertisers that look at this data are able to make a better inference about the impact of online tracked activity because of how many non-conversion Actions are taking place on the conversion tag. If the number is very small then the Advertiser-side analyst can be assured that the majority of sales are being tracked and attributed to paid-for activity

online. If the number of Actions is very high then it becomes tougher to identify the sale attribution for untracked channels and may indicate a need to extend tracking to other channels or marketing activity.

Conversion-based extended data at the Cookie Level

Cookie level reporting is also one of the few places where extended data parameters from the conversion tags can be captured. This would include the OrderID, Quantity of items purchased, and revenue generated, as well as customised parameters such as the names of the products purchased or their properties. These properties are presented on each row of the cookie-level data within additional columns featured after the conversion tag name and type, where a cookieID has performed a conversion activity.

Universal Pagecodes Mapping Table

Custom 1	Custom 2	Custom 3	Custom 4
(Location)	(Destination)	(Mode of transport)	(Seats)

[illegible]

Cookie Level Conversion Data Report

CookieID

Conversion
or Action

Figure 5.5.3 The Conversion Cookie Level data contains the extended data parameters passed through by the Advertiser script. Apart from the default parameters such as OrderID and Revenue, this includes the custom parameters. It is

worth the analyst consulting the trafficker to obtain the campaign setup information regarding the conversion tags so that a mapping table can be created to interpret the extended data. In this example a travel Advertiser collects data about flight bookings. Each custom parameter contains more information about the booking.

To enrich the cookie level data further, an Analyst at the Advertiser or Agency side can marry the OrderID field between the third party ad serving data and the Advertiser POS or CRM. This is best achieved bringing the two data sets together in a Data Management Platform (a DMP - see Chapter 8). Bringing these data sets together allows an Advertiser to build a full map of specific customer behaviour, since the cookieID becomes enriched with data that might offer the chance for the analyst to perform a method of profiling.

For example, the OrderID could indicate from the data in the CRM that the Order belongs to an account ID which had undertaken purchase actions over the last few years to indicate a compact buying pattern in the lead up the Christmas holiday with the same Advertiser. Using the ad serving data, the cookieID may have undertaken a certain pattern of site visits displaying the Advertisers ads before purchase. The analyst could see if there was a trend between users of look-a-like profiles and decide whether or not to inform the planners and buyers to consider buying more inventory on such sites closer to the end of the year to attract the buying behaviour from these users and similar users. Alternatively creative messaging could be adapted using DCO to perform a targeting function across inventory that had already been purchased on such sites.




Figure 5.5.4 the extended data from the Cookie level on the conversion tag can be married with the CRM information from the Advertiser side inside a Data Management Platform (DMP). Such systems enable the data to be connected with a programmatic optimization (see Chapter 8) in order to either adjust creative messaging or targeting for channels like DCO, adjust bid prices on keywords for paid search or auctions and bids in the programmatic space through a DSP (see Chapter 8). Of course the married data is also hugely insightful for the marketer.

Impression and Click based token data at the Cookie Level

Publishers can pass information into the impression call or the click pathway of the ad by populating a token field in the ad tag code when the ad tag code is fired. This is executed by a Publisher created script that sits on the Publisher page. Any extra information passed through in either of these dynamic token fields can be used to populate additional columns in cookie level data. Both of these fields are highly favoured by marketers that want to get a better idea about exactly where the ad appeared on the page. Since the Publisher ad ops team is the one implementing the ad, even page scraping the domain or page level URL where the ad was displayed is not enough information to convey where about on the page the ad sat. These fields can be populated

by placement position information such as 'position 4' so that when analysing the cookie-level data, the analyst can determine if ad performance or verification are affected by the position of the ad spot.

These advanced strategies take careful planning analysis. Some Agencies or marketing teams hire Data Analysts who process such data sets full time, for the purposes of campaign optimization. Analytics come with an important cost consideration when purchasing the media in the first place, so the full budget should not be blown on ad space but also on interpreting and presenting the results. Even bolting on other technologies like third party data providers (see Chapter 8) instead of physical analysts to get this data insight, can add significant cost and therefore affect the bottom line.

What is missing from ad server reporting?

Ad server reporting is absent one key component in both aggregated and granular forms. The efforts and the controls that the trafficker implements do not always have those setting fed down to reporting. This can mean that if the trafficker changes a conversion window, a rotation setting or the definition of an audience to omit a certain geography, these "optimizations" are not recorded in reporting. If the report displays a sudden spike in activity, the cause may be down to the optimization but only the campaign management section of the ad server interface will retain a record of the existing settings and may not event have kept an accurate log of changes to those settings.

It falls to the trafficker to maintain an accurate log of both campaign setup and any subsequent changes to the settings, so that reporting can be coloured with those optimizations that were responsible for the analytics output.

The next section focusses on optimization, a responsibility for the trafficker to undertake, in response to the ad server reporting information in order to improve existing campaigns, utilising ad server settings to do so.

Chapter 6

Optimization

Chapter 6

Optimization

- Questioning reporting data
 1. Where did the report come from?
 2. What kind of report is it?
 3. Do the numbers look right? Does the data add up?
 4. What time was the report pulled, what time zone is the account in and are these the correct date ranges?
 5. Are these the correct metrics, conversion tags and hierarchy entities?
 6. Do the names of the entities in the report (including the account name) match those in the campaign management interface?
 7. What is the methodology behind these metrics?
 8. Is the shown granularity providing enough or too much information?
 9. Is there enough data in the report or were the traffic volumes large enough to gauge decent statistical accuracy?
 10. Is there a chance that the data in the report be adjusted again later on?
 11. Are all the rows and columns displaying?
 12. Was the campaign setup in way that might affect a governed outcome?
 13. Did the ad server encounter any kind of counting fault?
- What is Optimization?
- Optimizing to Campaign Objectives
 - Optimizing towards Delivery Goals (ideal for Brand Advertisers)
 - Unique Reach
 - Target Audience

- Verification (quality safety) goals and impact
- Optimizations towards Performance Goals (ideal for Direct Response Advertisers)
 - Engagement and Interaction
 - Total Clicks
 - Total Conversions or E-Conversions
 - ROI (& extended data) goals
- Optimization Sample
- Recording changes in Optimization strategy

An analysis of the ad serving data early and often in a campaign provides an invaluable feedback mechanism to the Advertiser or Agency. By either making changes at a creative level or making changes to the planning and buying strategies, the outcome of the overall campaign can be influenced before its final impression has been delivered. Reporting therefore can act as the safety net to ensuring a campaign ends meeting the campaign goals or it can be used to take lessons from one campaign to apply to the next. The application of this learning to a live campaign is referred to as either a Media or a Creative optimization, depending on whether the media planner/buyer or the trafficker takes action.

The following section explains how to query the reporting data, before explaining how to understand it, to take positive action on it.

Questioning Reporting Data

Once the relevant report has been run and exported from the third party ad serving reporting interface. The analyst needs to make sense of the data. The following questions add much needed context to the data and ensure that it is of sufficient quality to make optimization decisions from. The analyst should know the answer to these questions before proceeding. If the answers are not available from the report or the help section

within the interface itself, the analyst can usually contact the support team at the third party ad serving provider.

1. Where did the report come from?

Having a colleague supply the data is not enough of a safeguard to ensure the numbers have not been tampered with by human hands. This is why third party ad server reporting interfaces grant access to a very wide range of users including the Publishers and the Creative Agencies, so that they can ensure the data is genuine.

Pulling a campaign directly from the third party ad serving reporting interface or securely receiving the data from the third party ad serving provider's delivery system (email or FTP) is the only way to ensure the validity of the source of the data.

2. What kind of report is it?

The report being pulled from the interface may not contain the right metrics to answer questions about the delivery or performance of the campaign. Pulling a report and running it, only to find it contains the wrong data can be frustrating and very time consuming.

Check the report description and title before running or downloading it. If it is not clear if it is the right report, ad servers will have supporting documentation or help sections to define the metrics it contains. In addition it is possible to display a preview of the report template without any data in it to see if it looks like it might be the right one.

3. Do the numbers look right? Does the data add up?

It pays to start with a hypothesis before looking at a report. This sets the analysts' expectations about the data that will

be revealed in the report. In addition it is worth tallying up the numbers manually to check that the total columns add up. If the numbers are wildly off, something might be wrong with the report, the data or the campaign.

If the data does not look right, pull the report again to see if the data looks any better. A single glitch or bug in the reporting interface for a temporary period of time could be affecting the way the numbers are outputted. If the data does not appear as expected, the trafficker can go back and test the ads, retargeting tags and conversion tags with *Fiddler*, checking that they right calls are being made to the server. These same pieces of code can also be tested in the live environment of the campaign to see if the code is implemented correctly or if the tags are being affected by other code on the page. If everything is performing correctly, read the rest of the questions in section and perform those checks before consulting the support team at the third party ad serving provider.

For reports that contain conversion data, check the conversion window settings and attribution settings for conversions. These may be causing the total numbers of conversions to be higher or lower than expected.

4. What time was the report pulled, what time zone is the account in and are these the correct date ranges?

Data processing occurs at different times and in different time zones. Processing can also be affected by the time zone that the Advertiser, Agency or Publisher account is set to. Pulling a report in a different time zone to the time zone on the account can mean that the latest data may not be available yet. In addition to this, it is common to pull a report with a date range that did not match the dates that the campaign was live.

To resolve this the trafficker can check the account settings in the campaign management interface. In addition the analyst can check that the date ranges in the report, match the ones required. It may sound obvious but a sizeable amount of queries about data are down to incorrectly matching date ranges.

5. Are these the correct metrics, conversion tags and hierarchy entities?

In haste it is possible to pull reports with the wrong metrics inside or looking at conversions for conversion tags that are not on the pages the analyst is looking to analyse.

When creating any naming conventions these should be shared with all parties so that it is clear which names match with what. Also analysts should be sure to understand ad server hierarchy levels, for new beginners it can be common to confuse a placement with a package for instance.

6. Do the names of the entities in the report (including the account name) match those in the campaign management interface?

The analyst may be working from information which is not true, being asked to pull conversion tag results for completely the wrong conversion tag.

The analyst can flip between the campaign management interface and the reporting interface to ensure that the right entities are being selected. The analyst should be wary that if the Advertiser does not use an integrated tag management system it is not possible to confirm which page conversion tags are implemented on. Analysts that suspect faults in conversion data should ask traffickers to

check that conversion tags have been implemented correctly.

7. **What is the methodology behind these metrics?**

The definition of a metric may provide an insight into why it is low, high or zero. For instance the definition of an impression according to IAB is to count when the server receives the ad request. An impression elsewhere might count when the ad is displayed on the page. Comparing reports of the same activity but different methodologies is not recommended, the data will never be the same. Also, not all counting systems can count things at exactly the same point in time. Usually one system is counting things before another which at ad serving scale can result in different figures when systems are compared.

To resolve this, the analyst should check to see if a known and accepted discrepancy exists between a comparative report. The IAB confirm that an accepted level of industry discrepancy for reports that track the same things in different systems (using the same methodology) should not be over or below ten per cent. There are three types of possible reporting discrepancy.

- a. Between data from the same source (reports pulled from the same ad server)
- b. Between data comparing third party ad server with Publisher ad server
- c. Between data comparing the third party ad server with another technology provider such as a 4th party or third party ad server

Although an ad serving provider can commit to investigating reporting discrepancies, type b. is much more challenging to resolve, since the two types of system report in very different ways and it is in the best interest of the Publisher

to report inflated figures. Instead Publishers can be granted access to the third party ad server reporting interface.

8. Is the shown granularity providing enough or too much information?

The most granular kinds of reporting are difficult to action without spending more time or resource on trying to extract the right insights. On the flip side, if a single ad tag had been provided to a Publisher to run on all page types of their website, reports could not give an indication of performance by site section for instance.

Getting to know which reports will display what level of data granularity will make the process of analysing and pulling reports less frustrating. Displaying a preview of the report template before running it will give a good indication of the level of granularity that will be available. It also helps to have a strong understanding of the definitions of the various levels of reporting granularity. If reporting does not seem to yield granular enough data, analyse the cookie level data instead. By comparing the domain URL captured from the page level where the ad was delivered and amalgamate up the results by subdomain, the analyst can get section summaries without having to traffic more ad tags. To save costs on analytics however, the trafficker should create more tags, replicating the same ad assignment to placement setup with section names, for the Publisher ad ops team to implement.

9. Is there enough data in the report or were the traffic volumes large enough to gauge decent statistical accuracy?

When dealing with very big numbers, very small numbers in reporting become close to insignificant. If a campaign was delivering one hundred million impressions, a placement or

a single ad version that had incremented a thousand impressions alongside versions that had incremented millions of impression, is a version that may just as well had incremented none at all.

The trafficker and campaign planning team should have accounted for distribution of data across the setup. This is relevant to the creation of audience segments and versions for DCO. On such a scale, wait for ad server entities to increment larger amounts of data before conducting an analysis. This is why it is recommended for the analyst to wait a day or two before pulling reports for the first time when a campaign goes live, ensuring that the data displayed are more than just the Publisher running a few tests on the ads while implementing them. If the data continues to be too small the trafficker can take the execution of the campaign setup altogether or simplify the experiment.

10. Is there a chance that the data in the report be adjusted again later on?

The analyst will want to conduct an analysis across the final results rather than have to undergo two rounds of analysis on almost the same data set because the numbers change. In particular, numbers in ad server reporting sometimes change for conversions and for Search channel cost data. Conversions can change when the ad server conducts several rounds of data processing owing to complex attribution modelling and de-duplication. The *Atlas* ad server used to run one round of conversion processing, delivery a near-accurate conversion figure to get early results available in reporting, before conducting a more thorough sweep and delivery within twenty four hours. Also *Google Adwords* sometimes recalculates search engine cost data over the course of a week, which is presented in ad server reporting.

In all instances check with the support team at the ad serving provider that the reporting numbers displayed will not change and get confirmation about timings and conditions on which the data can change. For conversion data, the analyst can wait until conversion processing is complete before pulling reports for the more accurate numbers. For search cost data, the analyst can pull reports daily for the previous week look back and compare those numbers to ensure cost data has remained fixed and is not changing.

11. Are all the rows and columns displaying?

Sometimes, particularly with large Advertisers, old campaigns can get archived by the ad server to speed up the time it takes to run big reports. This is also sometimes true of rows or entities that display no data or zeroes. This is useful to speeding up the performance of the reporting interface, but is an incomplete view for an analyst that may want to take old or low figures into account when analysing the data.

If a report is pulled and rows or columns expected seem to be missing, the analyst can run the report again, ensuring that the filters included are not filtering out certain metrics, entities or forms of data. Also it is worth the trafficker checking the account settings to see if and when the ad server is archiving any of the data. Where reports cannot be retrieved back to a certain date, the analyst can check to see how long the ad server keeps data stored for by asking the support team at the ad serving provider. In addition there is merit in asking the same support team to check that none of the old data is being archived.

12. Was the campaign setup in way that might affect a governed outcome?

Careful planning is essential in the campaign setup process. If the trafficker and the creative developer are less experienced, they may not have thought about the report outcome as a result of the setup. Creating too few tags for instance won't give enough granularity to provide the most insightful results. In addition there might be high numbers associated with more negative metrics such as incomplete video plays for Instream or no click data on a high reach Display execution.

Before starting any Rich Media campaign, the trafficker should have a kick off call with the creative developer and Publisher, making sure interactions are in the right place. Following this, the ad implementation should be thoroughly checked during the ad QA process. Once the campaign is live, verification reporting will report errors in targeting on the Publisher side such as ineffective brand safety. By checking delivery reports early and often, the analyst can avoid results that demonstrate an incorrect implementation by having the trafficker work with Publishers fix incorrect implementations. To reach certain goals, a campaign should be correctly optimized; this is covered in full detail at the end of the section.

13. Did the ad server encounter any kind of counting fault?

If the data reports significantly lower or higher volumes for delivery and performance there is a possibility there could be a counting fault in the ad server itself. This is a very rare occurrence and should only be considered as a concern as a last resort.

After all tags have been double checked and there still appears to be bad data in the report, the analyst can

contact the support team at the third party ad serving provider. This team will be able to conduct an investigation. In the rare occurrence of a fault, data can usually be re-processed and re-published to the correct time periods in ad server reporting. Depending on the severity of the issue depends on how long this fix will take. When considering using an ad server for the first time, the provider can supply information about data security and data integrity. The Advertiser or Agency can expect certificates of accreditation from well-known internet and hardware security institutions who vet the ad server to ensure that data faults and server downtime are virtually zero.

Once the analyst has established that the data is an accurate reflection of the live or completed campaign, the process of optimization can begin.

What is Optimization?

Optimization is utilising ad serving data to make an improved change in the campaign setup. This change can either occur in the ad serving campaign management screen or adjustments can be undertaken by the planners and buyers, changing which Publishers and audiences the ads still to run would run against. These are referred to as Creative Optimizations (in the ad server interface) or Media Optimizations (adjustments to the media plan outside the ad server).

Chapter 3 - DCO explained how the ad server can make use of the ad serving data without anyone needing to pull reports. The ad server would make use of the data in real time to adjust the creative optimization of the ads using "auto-optimization". The remainder of this section will look instead at the alternative, manual optimization. This means exporting reports from the reporting interface and having an analyst look at the data and

instruct the trafficker or the media planner to make changes to the campaign to improve performance or delivery. In this sense optimization might allow an Advertiser to reach a campaign goal faster or exceed it, ideally without increasing costs, but reducing them in order to do so.

The most stringent of Advertisers will insist that optimizations occur as often as possible. This means implementing an optimization strategy and then waiting for the adjustments to take effect before optimizing again. It is for this reason that having real-time time available data is not so useful for real-time optimization on small-scale campaigns. Optimizations will usually need a minimum of a day or twos worth of data to clarify if the adjustments have worked and daily aggregated reporting from the ad server is usually sufficient for this need.

Optimizing to Campaign Objectives

Optimization means making the media budget go further. First, the analyst can pull cost reporting from the reporting interface to see where the majority of the media spend is going, then look more deeply at the data to see if those Publishers or the Creative is driving ad spend to those goals effectively or having a strong contribution to it. If goals are not effectively being met (or to improve things further) the analyst can consider having traffickers and planners make some the following adjustments as optimizations, depending on the campaign goal.

Optimizing towards Delivery Goals (ideal for Brand Advertisers)

If an ad can be seen by the most relevant audience it can have the most impact, even if the Advertiser has not set any performance goals. By ensuring ads are delivered to the most quality, relevant pages, at the right reach and frequency and to the target audience, there is a better chance to create awareness with those users that encounter the ads.

Unique Reach

The optimization of Unique reach is a goal to find an untouched audience or recognising how much exposure a unique audience requires to cement the impact the ad can have on those users. It is advised to first set a baseline to establish how many unique eyeballs or users the campaign is reaching. To do this, pull a reach report by Publisher and by Channel. Some channels and Publishers may be more effective than others at reaching the intended audience. This data can then be combined with frequency information to show the levels of frequency that each Publisher and Channel is incrementing on each unique user. Combining this with cost reporting will show what it is costing to deliver at the baseline.

The next step is to consider the importance of reaching a unique audience. Usually the goal is harbouring a goal of performance to reach a certain number of sales. If this is a goal that can be reached online then the analyst can consider optimizing the campaign to conversions as a primary goal over unique reach. By analysing the attribution funnel in cookie level data, the analyst might reveal that might show that frequency is more effective than unique reach at driving that conversion goal, so frequency can be adjusted to match this need either on the third party ad server or with the Publisher ad ops team directly. Sometimes the goal for unique reach does not tie back to online conversions but instead might be designed to his an effective "GRP" goal. Here an Advertiser will combine the online reach data with reach data for other media channels such as TV to calculate, project and forecast sales offline.

Unique Reach - Creative Optimization

- To perform a creative optimization on unique reach (show a different message to new eyeballs) requires the use of retargeting (See Chapter 7) such that old users are labelled

and fall into a defined audience where new users are a separate audience.

- If retargeting is combined with DCO it will not cost the Advertiser much more to create personalised messaging for the retargeting audience, improving the campaign ROI.

Unique Reach - Media Optimization

- The unique site overlap report will show which sites are offering the best unique reach, an analyst running this report can have planners and buyers adjust their buying strategy accordingly.
- A reach and frequency report will reveal the effective frequency for sites, placements and ads so that the frequency cap can be adjusted at the Publisher ad server. Less budget could effectively be transitioned away from frequency to hit a higher unique reach goal.
- If a reach goal needs to be very high, it may be more effective to pay Publishers based on alternative payment models such as sponsorship or to utilise a channel like Affiliate.
- A device report will show high delivery to mobile devices that do not support third party cookies. These devices will hinder a unique reach goal since it would not be possible to determine if the same users are receiving ads cross-device.
- Analysts can also conduct an analysis on conversion tags for total "actions" and compare this data to site side analytics reports on referral information. If high volumes of traffic are reaching the Advertiser site without being tracked by the third party ad server and these users are arriving to the website from Publisher sites directly,

there is a gap of user traffic that the ad server is not capturing. The analyst can report these gaps to the planners and buyers to reach out to these Publishers to implement third party tracking and delivery. Without this in place, true optimization to unique reach goals becomes much more challenging.

Target audience

Use content classification verification reporting to establish the baseline and distribution of the audience on a non-specific audience buy across multiple publishers. The affinity between audiences and content or audience verification reporting specifically will reveal what audience groups are seeing ads. Once it is clear which geographies and inferred demographics the users belong to, the optimization can begin.

Target Audience - Creative Optimization

- Where higher cost ad formats such as Rich Media are delivering to an ineffective audience, the trafficker can add audience targeting to the campaign setup in the third party ad server interface.
- Fluctuations of user traffic for various demographics are known to map to times of the day. Rather than pay for additional audience targeting data, the trafficker can utilise ad controls to show different creative messages at different times of the day.
- To lower creative costs, the Advertiser can utilise DCO and retargeting for personalisation as well as creative message sequencing.

Target Audience - Media Optimization

- Since Verification reporting for content and audience verification, will reveal the audience types that Publishers provide, planners and buyers can use this

data to spend less with Publishers that do not attract the target audience. This saves money by circumventing the need to increase spend on a guaranteed audience targeting with that Publisher.

- If planners and buyers invest in inventory with a price based guarantee against a target audience goal, such as geography, there is an optimization opportunity. When the activity goes live, the analyst can run a verification geo report to catch Publishers that fail to setup the geo targeting. The remuneration of the failure can be spent on audience targeting with a Publisher that does not breach the audience targeting conditions.

Verification (quality safety) goals and impact

To deliver more quality impressions and user viewability, the analyst begins by comparing campaign verification and viewability results with the benchmark report for the same Advertiser vertical as the Advertiser in question. Alternatively it makes sense to look at verification reporting from previous campaigns with the same Advertiser across the same Publishers. This will set a baseline from which the Publishers and Creative in use must meet.

Verification - Creative Optimization

- Viewability reporting or Rich Media reporting can be used to reveal the average view duration that users look at ads for. The trafficker can use this information to remove ads from rotation that have low view duration. The same optimization can be made with Instream reporting.

Verification - Media Optimization

- An analyst can use the brand safety and domain URL features of verification reporting to produce a blacklist

of sites that breach the brand safety requirements. By issuing this blacklist to the Publishers, optimizes the inventory by ensuring ads are not shown again the blacklist in future. This proactively negates the risk of a high cost associated to bad PR around a failure of brand safety. This cost can add hindrance to the cost of running the campaign in the first place.

- Pre-emptive verification utilises a similar strategy. The unseen benefit is that inventory which adheres to verification requirements is considered a higher quality and lower quality inventory could be bought at a lower cost to satisfy a secondary reach goal.
- A buyer could introduce a new cost model to pay on viewable impressions only. This penalises poor viewability in the marketplace and means Publishers may make more of an effort to optimize media to appear in view. The in-view guarantee may inflate media costs which the buyer must consider worthy of the in-view goal before proposing the cost model adaption.

Optimizations towards Performance Goals (ideal for Direct Response Advertisers)

Determining the performance of the ads is seen by most Advertisers as the most important consideration when running digital campaigns. Any trends in reporting which can be shown to boost performance are worth carefully monitoring and applying to a wider audience reach. It is important also for the analyst to consider not just optimizing to one performance goal. High levels of engagement might diminish clicks or conversions for instance, so gathering a wide set of data before presenting the argument for the optimization to the trafficker and marketing team is a best practice. As with Delivery goals, storing old performance information or looking at publicly available data for the purposes of benchmarking allows expectations about performance to be set

from the beginning of the campaign. Performance goals exist for engagement, clicks, number of conversions and value of conversions as an ROI.

Engagement and Interaction

Interaction and engagement is only going to be available on Rich Media and video ads. Although the cost to serve formats in these 'channels' will be higher, engagement can be a valuable indication of performance if creative and media fail to instigate users to click, convert and buy from ads.

Brand Advertisers that rarely look at performance would do well to monitor engagement levels in reporting. At the least, they can provide an indication of awareness.

To establish a baseline for engagement and interaction as a success criteria, the analyst can begin by checking the benchmark and performance reporting for engagement on previous campaigns. At a minimum, format performance for engagement by country and Advertiser vertical will set an expectation before optimization begins. The goal is then to beat the benchmark or improve on the previous campaign.

Engagement - Creative Optimization

- Sometimes creative developers and designers recycle old creative templates to save time but user engagement has an affinity towards novelty and difference. The Creative Agency can make the ad more appealing to the user by using the latest available components in the ad serving provider supplied component pack for Flash. Inspiration for new executions can be attained from ad serving provider showcases. Some providers also offer in-person support to inspire innovation based on the available technology. More innovative executions will boost engagement levels.

- Creative designers can now utilise HTML5 to deliver executions with better opportunity for engagement through motion, touch and device-specific features. The Advertiser may wish to consider paying the Creative Agency to deliver a new execution, if the campaign is already running, if the campaign is not performing well on engagement across the board.
- Instream Video, runs at a higher media cost and cost to serve but is shown to provide significant opportunity for user engagement. In particular the analyst may have to consider switching from VAST formats to VPAID or DCO for Instream. These formats are emerging and will attract higher levels of engagement due to their novelty.
- The analyst can look at ad server reporting by format cross-Publisher. The reports may indicate that some formats attract better engagement than others. The analyst can instruct the trafficker to drop low engagement formats from the placements or up-weight those with higher levels of engagement. It is possible that if the frequency and reach levels of badly performing formats were adjusted that they might improve to even higher levels than those ads that appear to have high engagement at low frequency rates.
- The analyst should also consider looking at cookie level data to extract information about what ads and formats may have influenced levels of engagement in events on the pathway to a specific custom interaction.
- In addition optimizing to a refined audience targeting may reveal an affinity between audience types and engagement levels to present the rich media ads to those audiences that are more inclined to engage and display ads to the rest.

Engagement - Media Optimization

- Analysts may wish to look closely at engagement reports and interaction metrics to see which Networks, Publishers, Sites, Site sections, Placements and whole channels are driving the optimum interactions. The analyst can advise the media planners and buyers to adjust the media plan according the results to improve the campaign performance.
- Traffickers can place retargeting tags (see Chapter 7) on the custom interactions within the creative and then programmatic buying (see Chapter 8) could be used to seek out those same users again by buying media that they are actively on. By extending further media cost to those users above all others, the optimization is being undertaken on engagement.
- By introducing a cost model based on engagement, such as cost per engagement or pay out on a video duration metric, the campaign will self-optimize towards this goal.

Total Clicks

The analyst could begin optimizing to clicks by establishing a baseline: checking the benchmark and performance reporting for clicks on previous campaigns. At a minimum, format performance for clicks by country and Advertiser vertical will set an expectation before optimization begins. Previous effective click cost data (based on a CPM) will also provide an indication as to the expected number of clicks based on the spend alone.

Total Clicks - Creative Optimization

- Standard reporting will indicate which ads and creative drove the most clicks from within the placements. The trafficker can adjust the ad controls such as weighting on

placements so that ads and creative attracting more clicks are shown more often. If different creative is being shown to different Publisher sites, consider swapping the creative over if creative of the same dimensions with different messaging is performing better elsewhere as long as the Publisher can be ruled out as a reason for the improvement. Collating all historical reporting together can help form a benchmark to be clear if the Publisher is responsible for the improvement. Since the adjustment to the creative happens in the ad server, the live campaign will not be affected by the change.

- If creative messaging is not standardised across all ad sizes and one message is found to be responsible for improved performance, consider having the Creative Agency build additional ads with the optimum message. The trafficker can then put these into rotation. To cut creative costs, consider using DCO from the outset to take advantage of dynamic text fields.
- In this process the analyst can analyse the user pathway to the click events. A pathway to conversion report may indicate that clicks are influenced by other ads and creative messages.
- Some ad servers filter multiple clicks on the same ad. If multiple clicks by a single user can still be classed as a contribution to the goal, consider converting the display ad to rich media and counting additional clicks using custom interactions.

Total Clicks - Media Optimization

- Pulling campaign reports will indicate which sites, sections, Publishers, Networks, channels and formats provided the most clicks. Planners and buyers can consider revising the proposed media spend, distributing

the remaining budget to better performers where possible. This may include excluding some Publishers from the remainder of the campaign.

- Before removing anything from the plan, the analyst can look at pathway to conversion and cross channel reporting to ensure that any changes to the media plan will not inadvertently affect click performance elsewhere.
- Pathway to conversion may also provide stronger case to demonstrate that some Publishers are adding no value across the board.
- For search specifically planners and buyers could pay more for the keywords that drive the most clicks. Keyword path analysis can play a role in showing which keywords and combinations of keywords are driving clicks across all tracked channels.
- The analyst can use cost data to determine how much it is costing to get a click, even if none of the activity is running on a CPC model by calculating based on the impression to click ratio (CTR). Use this to calculate an effective CPC. Propose paying Publishers on a CPC is an unlikely ask but it will ensure that the media is totally geared towards click performance. Use the effective CPC to convince the Publisher to lower the cost of the media so that the impression costs will yield the click costs based on the performance to date. A CPC model might be possible across affiliate activity.
- Channel reporting may reveal that some channels have a greater impression to click ratio such as email. So long as the reach volumes are high enough to reach the click goal with the available inventory, use the ad serving data

to demonstrate that planners can spend more across those channels.

Total Conversions or E-Conversions

The analyst can begin optimizing to conversions by running conversion reports by format as well as by Publisher. This will indicate where and when conversions are occurring. Establishing a baseline for conversions as a success criteria, can be more challenging for conversions, as the amalgamated conversion data in the benchmark might include conversion events that are not comparable to an Advertiser's goal. For this reason, previous campaign data is more valuable to set expectations.

Audience quality plays an important role in conversion optimization. This does not always mean making an adjustment to planned delivery by Publisher. It can mean acquiring more data to provide better audience intelligence. For instance, those users with insurance policies up for renewal may have made this information available to a data collection company. Marketers could use the marketing budget to borrow the data to establish which cookieIDs are more valuable to the conversion goal. An advanced data management strategy could achieve this using a DMP and be followed by a creative or media optimization. In addition it is important to decide what conversions are the established goal. Selecting view based conversions or click based conversions come with their own considerations. Also, with conversion optimization, where click-to-land is priority for the Advertiser, click based conversions would take priority and the attribution settings on the conversion tag can be setup to meet this need.

Total Conversions - Creative Optimization

- Just as with clicks, standard ad server reporting will show which ads and creative drove the most conversions. The trafficker can adjust the ad controls accordingly so that

these ads are shown more frequently. As with clicks and engagements, the analyst can consider how other ads and creative in the same campaign contributed to the conversion before making any changes.

- Conversions are also affected by the Advertiser site, so analysts can also focus on the clickthrough URL to establish if conversions were more likely if the user arrived on a particular landing page.
- Advanced Time to convert reporting will reveal if the conversion window has been set too liberally or not. Having the trafficker make the adjustment to the windows on the conversion tag based on the analysis will ensure conversions are not being missed.
- The trafficker could also carve off some of the audience to see if A/B testing exposes that one audience performs better on conversion events than another. The traffickers can then work with the analyst to optimize to this audience.

Total Conversions - Media Optimization

- Once again, the analyst can look at which sites, sections, Publishers, Networks and channels provided the most conversions and consider removing those that are not driving as many conversions or any at all from the media plan.
- For e-conversions, it will be essential for the analyst to look further up the conversion funnel in cookie level data at the pathway to conversion. Although engagement mapping will take into consideration the conversion influence by other factors, the analyst should keep checking reporting to ensure the most effective

attribution model is in place for the optimum representation of e-conversions.

- With regards to cost model in display, the media planners and buyers could agree to pay Publishers based on an agreed attribution model as a CPA. This will ensure the Advertiser's unique conditions for conversion optimization are being met in the buying cycle.
- For the search channel, the media could be optimized by the trafficker by paying more for the keywords that drive the most conversions and bringing in P2C data to enhance the decision.

ROI (& extended data) goals

The ROI is revealed in the conversion tag extended data, so the analyst will want to look at cookie level data on the conversion tag. Where the extended parameters have been populated the cookie and the pathway to conversion are the key to ROI.

Comparing costs for media and technology with the revenue figure (real costs for the purchased item) will reveal the ROI.

ROI Creative Optimization

- The analyst can determine if there is a link between creative, format and the final revenue.
- The trafficker could implement a DCO campaign to lower creative costs across the board, improving ROI whatever the final revenue figure turns out to be.

ROI Media Optimization

- The analyst could build profiling around those cookies that spend more money and target users which have similar behaviours with a media and creative strategy that worked to get the profiled users to buy.

- The analyst could establish a baseline for the audience based on content verification reporting and by Publisher to ensure that planners and buyers adjust the media plan to buy across more affluent audiences.
- As with other forms of media optimization, the planners and buyers could choose to pay Publishers based on an ROI model. This would be articular effective across performance channels such as affiliates.

Optimization Sample

The audience size for any optimization (both manual and automatic) will also dictate the success for the optimization strategy. A very small audience size does not provide enough data to give a true indication of performance when scaled up (because in a smaller group some users do not behave according to the general population but can swing favour if the consideration set is low).

Recording changes in Optimization strategy

Successful strategies can be kept on record and potentially reapplied to campaigns under the same or similar conditions later on. Traffickers can record where and when changes to a setup are made, what data drove the choice to make the optimization, and the outcome of those optimizations so that those that yield positive results can be re-applied later on or in a future campaign.

We have looked generally at how reporting data can be applied to campaigns for the purposes of optimization but one form of optimization has proven to be consistently effective at reaching some goals. In particular, it is cyclical nature of auto optimization that has driven high adoption for behavioural targeting, the retargeting to users that have encountered the Advertiser before.

Chapter 7

Retargeting and Privacy

Chapter 7 – Retargeting and Privacy

- Retargeting Tags
- *Figure 7.1*
- Cookie Pools
- Retargeting Windows
- Retargeting for Creative Optimization
 - *Figure 7.2*
 - Dynamic Retargeting
 - *Figure 7.3*
 - Carousel Ads
 - *Figure 7.4*
- Privacy and Behavioural targeting
 - The Ad Choices Program
 - *Figure 7.5*
 - Do Not Track
- Retargeting for Media Optimization

Retargeting is a way to reach leads and former customers over the noise of other ads online and is considered to be the most evolved and effective optimization strategy. Using retargeting tags allows the trafficker to plan how creative messaging and media buying will automatically be adjusted if the ad server has encountered or not encountered a user before. This Chapter will explain the workings of retargeting and how it adds an extra dimension to DCO enabling a programmatic media buying strategy. In this sense, retargeting can be classed as an automatic optimization as oppose to a manual optimization.

Retargeting Tags

For retargeting to work, the trafficker needs to generate a retargeting tag from the campaign management interface of the third party ad server. Retargeting tags behave in a very similar

fashion to conversion tags. These are small pieces of code that, like the conversion tags, are sent to the site management team at the Advertiser site to implement the tags into places on the site which would benefit from capturing user behaviour. Such locations include product pages, landing pages, pages in the booking or payment process and the "thank you" or confirmation page. Strategically placed, these tags can be set to trigger when the user behaviour is performed. Examples include when a user prematurely closes their browser window before completing a transaction, or demonstrating some interest in a product of a particular size and colour. The flexibility of these tags means that the strategy for the Advertiser or Agency can be complex; building groups of users that have exhibited a particular on-site behaviour.




Figure 7.1 Retargeting tags are flexible and can be added to multiple locations to build up groups of users that have performed similar actions or have arrived on similar pages.

Since retargeting tags can also be piggybacked on conversion tags they can capture some of the same extended data parameters that a sales conversion tag can capture. This allows for retargeting strategies that only show certain ads to users that, for instance, spent a certain amount of money on a purchase. Retargeting tags can also be placed inside ads themselves during the development and setup stages of the campaign. Such retargeting tags can load

when the ad loads or when the user triggers a certain interaction or behaviour inside the ad. Some Advertisers adopt a policy of ensuring that users should never be allowed to encounter an Advertiser without being labelled using retargeting. This method would ensure that future efforts to message this user would become more personalised. This is based on knowing that a strategy used in the past to target this user had failed to incite a conversion and that a different strategy should be adopted.

It is not uncommon for Publishers to restrict the use of retargeting tags inside ads on their sites and networks. Building such retargeting groups effectively diminishes the value of the audiences that the Publishers are attracting and trying to sell against. This was referenced in Chapter 2 as "Publisher Leakage". Having the creative developer and trafficker, checking Publisher specifications before implementing a retargeting strategy into ads directly, will ensure that no expectations are being breached once the ad tag code is trafficked.

Cookie Pools

Retargeting Tags can be implemented during a campaign for the purposes of optimization, but work best if placed before the campaign begins, to ensure that users can be targeted based on this information from the outset of the campaign. When a user encounters a retargeting tag their cookieID is placed into retargeting group or "cookie pool". This pool is a way of categorising that user so that they can be messaged later on.

Some ad servers allow Advertisers to see how many users are in a specific cookie pool. This allows the Advertiser to know if it is worth setting up a campaign or ad specifically targeting the group. Very small pools are not always worth considering because on a big scale the effort to target these users later on can cost more in man-hours to set up and manage, than the potential increase in earning by targeting those users again.

Retargeting Windows

The aim of retargeting is to show an ad to that same user at a later date with the hope that by seeing the ad, they will be contributing to the campaign goal. The opportunity to build up a pool of cookies and show that pool targeted ads at a later date is attractive to marketers but is limited by a time frame which is rarely adjustable. Most ad servers offer a 90 day window in which those users in the pool can be retargeted. After that point, the cookies are removed from the pool (but as long as cookies are always being added to the pool, it will not diminish entirely).

Retargeting for Creative Optimization

Retargeting strategies for large Advertisers are rarely simple affairs. If a user can perform a number of actions on an Advertiser site or if a web-based product catalogue is particularly large, then re-targeting rules in the ad server need to be structured and prioritised. Users that go to a product page and then purchase and pay via a gift voucher, present some interesting possibilities for the Advertiser later on. Should a user encounter an ad from the Advertiser later, the ad can be designed to use the retargeting information so that the creative does not contain a reference to that product again.

Within the ad server interface, a retargeting audience is treated as any other audience type. In Chapter 4 we explored how the ad server can be setup to consider if each user is a suitable match for an audience type before moving on to the next 'rule'. The trafficker has the option to setup ads on the placement level so that when a user encounters an ad, the corresponding creative can be shown. This option works most effectively with DCO. The trafficker can ensure that a corresponding version matches with the user behaviour that triggered the retargeting tag.

For instance if a user visits the basket page but never reaches the confirmation page, two retargeting tags can be setup, one on each page. The trafficker can setup a target audience in the on the

placement level that declares that if a user is not in the cookie pool for the confirmation page retargeting tag (but is in the one for the basket page) then the version displayed should have a message related to the user abandoning to shopping process. To make the creative message more relevant to the user, it could be adjusted to specifically reference the items that were not purchased. This requires a more complex form of retargeting called "Dynamic Retargeting".




Figure 7.2 Dynamic retargeting tags allow Advertiser to pass dynamic variables into the retargeting tag so that the variables can be used in the version or ad later on. In this example the SKU number (a unique product identifier) is passed in the retargeting tag as the user cookieID is placed into the retargeting pool. When the same user navigates to the Publisher page, the displayed ad or version is targeted based on the SKU number, featuring the same product again or displays an ad which corresponds to the product. The decision is made based on the rules setup for targeting by the trafficker.

Dynamic Retargeting

Retargeting tags have the ability to capture a unique identifier, an example of which would be a product ID. A script can be implemented by the Advertiser site management team, allowing the page that the retargeting

tag is sitting on, to pass a value into the tag. As a user arrives on the Advertisers page (such as an e-commerce page featuring a jacket) the retargeting tag on that page receives the product ID passed in via the script. The retargeting tag then it sends a call back to ad server to add the user to the retargeting pool specific to the dynamic retargeting tag while writing the specific variable to the user's cookie. This way, the user becomes linked to that variable and the ID can be recovered when the Advertiser comes to show that user an ad later on.

The user is now associated with all other users that encounter this same tag but the difference with a dynamic retargeting tag is the deployment of the tag. A dynamic retargeting tag can be deployed to all product pages on the Advertiser ecommerce site as a single piece of code. The unique identifier passed into the tag is stored against the user's third party cookie.

When an Advertiser comes to run DCO activity the ad serving interface allows the trafficker to request that the version of the ad to be displayed depends on the product ID fed into the selected dynamic retargeting tag. Having connected the versions to a product feed for mass versioning, where all the product IDs are displayed, the result is that the user is per shown a version displaying the same product again. When the user encounters a placement, the code checks the user's cookie, passing the dynamic variable through to the ad server. At the placement level, the rules setup by the trafficker check to see if the user is within the cookie pool for the dynamic retargeting tag. If they are, the ad server matches the dynamic variable from the user's cookie with a corresponding ad version and displays this version to the user.




Figure 7.3 Dynamic retargeting connected with DCO. The versions of the ads for the campaign are built out based on the Advertiser product XML. This same XML is used to populate the product pages of the Advertiser site. The same SKU identifier is used to match the page with the version. The user enters at Advertiser site at (1). The product SKU number is passed into the dynamic retargeting tag (R). The dynamic retargeting associates this value with the user's cookieID and passes the information into the ad server (2). The user then navigates to the Publisher page (3) the placement carries the users cookieID to the ad server (4) where the targeting rules see that the user's cookieID is associated with an SKU number in the retargeting pool. The version with the corresponding SKU number is displayed on the page (5).

To make things more interesting, the trafficker and planning team can plan ahead and choose for the dynamic text field in the version to display a cheaper offer than the price shown on the site for the product. The dynamic click

through URL within the DCO version can be used to carry a token to the Advertiser page that could be used to dynamically change the physical price of the jacket. Another strategy is not to display the same product, but to display a complimentary product as a version of the ad.

Carousel Ads

Dynamic creative can be taken a step further by not just capturing the last item a user was exposed to, but to store several items or variables in the user's cookie. This can correspond to a historical record of the user's journey on the Advertiser site, capturing all of the product ID's encountered on the way.

When setting up Carousel or Catalogue ads, the trafficker has the option to utilise this data to display all previous products seen, or only certain products inside the delivered ad version. These could just be the most expensive products or those with the highest available stock.




Figure 7.4 Carousel ads keep a record of the users dynamic variable engagements on the user cookieID. The user navigates the advertiser site (1) where the dynamic retargeting tag captures the SKU number for the product in question. The user then navigates to the second page where the same thing happens. The cookieID records against both variables (2). When the user navigates to the Publisher site (3) the ad server serves an ad which displays both previously seen products by looking at the variables on the matching cookieID (4) and bringing versions together as a new ad version displaying the carousel of products in the placement on the Publisher site.

Privacy and Behavioural targeting

Governmental and legal bodies the world over are in support of internet users having a right to control the information that is collected about them. Strictly speaking the third party ad server is collecting information only about the cookie and not about the

user themselves (the information is not personally identifiable "PII"). It is the identifiable data that users are keen to eradicate and the cookie is not storing any of that information. Despite this fact, legislation is emerging to protect users, so that any behavioural targeting information collected about them can be made anonymous (or pseudonymous) or removed from databases completely. This relates directly to the topic of retargeting, since opting out of retargeting will mean that ads no longer to appear to 'chase' users across to different pages, Publishers and channels.

The Ad Choices Program

The ad choices program is a system that allows users to opt out of behavioural targeting via a small icon in the corner of the ad. The icon can be implemented by the trafficker inside the third party ad server even during a campaign, and the icon and its functionality will be added to a visual layer on top of the creative. When clicked the user is taken to a page where they can choose to opt out of retargeting altogether or by each technology provider or network collecting user data.

Point solutions by *Evidon* and *Trustee* have emerged to provide more complex management facilitates of the user opt-out including statistics for Advertisers specifically. These services come at an additional cost.




Figure 7.5 The ad choices program requires that all behaviourally targeted ads carry an ad choices icon. When clicked, the icon provides the user with information about retargeting and allows the user to navigate to the ad choices cookie opt out page. On this page the user can select or deselect cookie level technologies such as third party ad servers to opt out of. If an opt out is chosen, the ad server is informed. In future when the same user navigates to a Publisher site, the ad server will check the cookieID first to see if the user has opted out of behavioural targeting before deciding which ad to display. Those users that have opted out will no longer see targeted ads.

Incidentally, opting out of these pools using the opt out methods provided by such programs, requires the dropping of another cookie on that users machine to remember their settings. Deleting of these cookies, removes the memory of their choice to opt out and so users must opt out again after the cookie is expired or deleted.

Do Not Track

To get round this browsers are attempting to implement a cookie-free catch-all for the opt-out or opt-in of cookies

which is called the DNT functionality ("Do Not Track"). Discussions between legal parties, browser operators, governmental bodies, regulatory bodies and politicians continue to wage, as consensus around opting in and opting out at different levels of informed consent is mismatched. What is important to note is that opting out of behavioural targeting via a cookie opt out, will not stop ads altogether, but rather stop the occurrence of targeted (and some would say relevant) advertising which seems counter-intuitive to the best interests of all advertising stakeholders including consumers.

DNT is a privacy standard that is sends through a flag to the ad server on the http call and if it is turned on; servers shouldn't set up tracking cookies for the specific http session. This means that no cookie is dropped on the device. If a user disables the setting, the publisher ad server calls the third party ad server to add a parameter to allow cookie dropping. If the user enables the setting, the third party is called to disable cookie dropping. This will also disable piggybacked technologies from placing tags.

More concerning for Agencies and Advertiser is that disabling cookies will eliminate ability to do attribution reporting and count unique audiences for the purposes of de-duplication. While giving the users a choice to opt out of behavioral targeting and cookies may seem alarming to some marketers, when given a choice, few users choose to disable cookies altogether, according to IAB research. However, some browsers are choosing to ignore consumer choice and are instead force the disabling of some cookies, in particular third party cookies that would allow for retargeting. This directly damages the ability for third party systems to operate effectively and fairly and should be addressed with concern by governmental bodies against the browser's ability to over-ride consumer choice.

Retargeting for Media Optimization

Retargeting technology can also be used for media optimization, literally buying audiences based on retargeting tag encounters. This is enabled using DSP and programmatic technology in the ad exchange space covered in full detail in the Chapter 8 - Programmatic. At a high level, the DSP technology providers offer their own retargeting tag technology. This tag can be piggybacked alongside any other retargeting tag such as inside an ad server container tag. The user encounters the retargeting tag from both providers and as such is added to the retargeting pool. Within the DSP, the retargeting pool can be used to identify potential impressions as they become available on the ad exchange. If there is a match, then the Advertiser or Agency can bid on the impression through the DSP. Once the impression is won, the ad server is called, whereupon the trafficker setup of the ad checks to see if the user is in a retargeting pool from a creative perspective. They will appear to be so as the user triggered both of the retargeting tags and the user will be shown the corresponding creative. Tech providers such as *Criteo* have made a very lucrative business out of this offering but tying together the creative and media optimization capability of retargeting to buy and optimize against retargeted users at enormous scale.

This chapter has quietly offered a sneak peek into the topic of the next Chapter - Programmatic. Which, largely due to the success of retargeting campaigns, is seeing phenomenal growth and represents a true evolution in display advertising. Here, smart decision engines begin to replace the manual work of optimization and buying strategies. This replaces long operational man hours by agency and marketing teams to date so is being welcomed as game-changing by industry stakeholders.

Chapter 8

Programmatic

Chapter 8 – Programmatic

- Advertising Exchanges
- Supply Side Platforms (SSPs) for the Publishers
- Demand Side Platforms (DSPs) for the Advertiser
- *Figure 8.1*
- Connecting the ad server to the DSP
- *Figure 8.2*
- Quality Inventory Sources
- The Second Highest Bid Auction
- Real Time Bidding (RTB)
- Cost Reporting in the ad server
- DSP Strategies
 - DSP Retargeting with DCO
 - *Figure 8.3*
 - DSP Keyword Targeting
 - Negative Retargeting
 - DSP delivery controls
 - Building Private Exchanges
 - *Figure 8.4*
 - Strategies based on Publisher Audience Descriptors
 - Cookie Level Data Providers
 - Data Management Platforms (DMP)
 - *Figure 8.5*
 - Page Level Data Providers
- Channel-Specific Complications in RTB
- DSP - ad server of the Future?

Programmatic describes a fleet of technologies with the vision to automate the planning and buying processes of digital campaign setup using the power of reporting and audience data. These technologies are still in their relative infancy and are evolving but the infrastructure they create is very real and rigid, founded on the principles of the financial exchanges. This Chapter will provide an overview to the different technology types that exist within this space and explain how ad serving 'plugs in' to the programmatic stack.

Advertising Exchanges

In the traditional model of buying ad space, the Advertiser or Media Agency would contact the Publisher directly and make a deal to purchase inventory for the proceeding ad campaign. This was followed by the emergence of sales houses or Networks that worked on behalf of some Publishers, to sell the space to Advertisers. Networks aggregated the inventory together so that they could benefit from a wider pool of available ad space for the purposes of delivery fulfilment. On the relationship front, Advertisers and Agencies had relationships with both Networks and the Publishers that sat outside of them.

Network and Publisher negotiations person to person are a slow, time intensive processes relying heavily on relationships and to an extent; industry politics. By having machines make the negotiations, the process of buying and selling inventory becomes much faster, more efficient and more cost effective. It is now possible for Publishers to load their available inventory into a marketplace called an "Advertising Exchange" and Advertisers and Agencies come into the marketplace to buy and bid on the available space.

There is not one major exchange, rather a few that have been setup and some exchanges are channel-specific so that exchanges for video inventory exist such as *Tubemogul* separate to exchanges for mobile inventory. There are also broader exchanges

that exist, attempting to be established as the main marketplace. Major Networks have transformed their businesses from Network into an Exchange. This has given rise to the *Microsoft Exchange*, *Google's DoubleClick Exchange*, *Right Media* (Yahoo's offering) and marketplaces from more traditional Networks such as *HiMedia*.

Supply Side Platforms (SSPs) for the Publishers

Supply Side Platforms (SSPs) have been developed to bring Publisher inventory into the marketplace as soon as it becomes available. SSPs take the available inventory from the Publisher and classify or label it in order to sell it. Since Advertisers are looking to buy media space based on a target audience, the SSP gathers this information from the Publisher as a database that maps pages and placements on the Publisher side to audience types and demographics. Major SSPs include *Google's AdMeld* and the *Rubicon Project*.

Demand Side Platforms (DSPs) for the Advertiser

Advertisers and Agencies also have a gateway into the marketplace. This gateway is another technology called a DSP (Demand Side Platform). Heavyweight players in the DSP environment include *Turn*, *Mediamath*, *Appnexus* and *Google's Invite Media* and there are more emerging offerings to Advertisers such as *DataXu* and *X+1*. These systems are operated by what are referred to in the industry as "trading specialists" and the environment of buying and selling in the ad exchanges is termed "trading". Trading specialists will typically work either directly on the Advertiser marketing team; at the Media Agency or for an independent outsource company. As explained in Chapter 2, these teams are referred to as "trading desks" and each trading desk may work to access the inventory in the exchange via one or more DSPs. Those that sit within the Agency are referred to as ATD's (Agency Trading Desks) and those that are outsourced are called ITD's (Independent Trading Desks). By working continuously in the programmatic environment, trading

teams can maintain a history of programmatic buying activity and cost information, much like the planners and buyers would do when buying digital media non-programmatically with the assistance of a planning tool.

Using a DSP is the overhead cost of accessing what is referred to as the Open Exchange. In other words, the DSP is the Advertiser's ticket into the marketplace from where they are free to browse and buy all the inventory on offer.

The DSP technology is very similar in its working to a Search Engine Marketing platform (SEM) such as *Google's Adwords*. The DSP operator has a record of the available media budget for the campaign and sets goals to buy target audiences within a certain period of time. The combination of these variables is called a "strategy" and the trader will setup a strategy and put it live to allow the DSP to enter the marketplace to seek out inventory that matches the variables. To set a strategy live requires the trading team to connect the variables to an ad so that if the correct conditions are met; 1) the Advertiser pays the Publisher through the marketplace 2) the ad is released from the DSP to the SSP, where 3) it is presented to the user on the Publisher page via the Publisher ad server.



Figure 8.1 Programmatic simplified. The open exchange is a marketplace where sellers present the available ad space via the Supply Side Platform and buyers examine and pay for the space via the Demand Side Platform. When the deal is met, the DSP

provides the ad to the SSP from the ad tag provided by the third party ad server.

Connecting the ad server to the DSP

To have this work requires that the trafficker, working in the third party ad server campaign management interface, set up a placement against a site entity which has the same name as the DSP or the trading desk. For example, if the media plan received by the trafficker from the Media Agency stated on one of the rows that an ad tag needed to be trafficked for the *GroupM* trading desk – *Xaxis*; the trafficker would setup the campaign within the ad server interface such that the site was *Xaxis*. The trafficker would then assign the right placement and ad to that site and traffic the ad tags to the *Xaxis* team. The *Xaxis* trading team would then load the ad tag into the DSP interface and the DSP would connect to the open exchange to find a suitable ad spot.



Figure 8.2 *The programmatic space fits into the buying and selling process between the Advertiser and the Publisher. The direction of the arrows in this image represents the direction of the ad and of the media budget. Greyed out is the traditional buying model: Media Agency or Advertiser buys the space directly from the Network or Publisher and reserves the inventory. The creative agency then loads the creative into the ad server, the tags are passed to the network and the ad is served at the right time. With programmatic in the picture the ad is only passed to the Publisher when the agreement or strategy is fulfilled and the process is managed by the trading desk.*

Quality Inventory Sources

DSPs will typically connect to multiple exchanges in order to fulfil a strategy but this is not a default setting and something an Advertiser should check when using a trading team or DSP for the first time. The various marketplaces carry inventory of differing quality.

The idea of the advertising exchange emerged to try and sell off all of the available ad space. Advertisers and Agencies would traditionally use the Publisher sales teams to find quality or premium inventory to display their ad within. They would be willing to pay a higher cost for the space. In such an environment, bottom-of-the-barrel or "remnant" ad inventory went unsold.

The ad exchange provides an environment to sell ad space based on the price that Advertisers are willing to buy it for to reach their goals. This has led to some exchanges carrying more valuable ad spots while some are happy to carry the remnant. The major debate of contention arises in this space where the definition of available quality inventory must be very concise and clear so that the goals are met and the right price is paid without giving either the Publisher or the Advertiser a bad deal.

The Second Highest Bid Auction

Programmatic can work by having an SSP present ad space, available in the future in the exchange. This inventory can be presented in multiple exchanges so as to increase the number of DSPs that can see it. Once the inventory is presented, the DSP strategy will identify the inventory as applicable to the defined variables and will make an offer for the ad space. Other DSPs are invited to do the same on behalf of other Advertisers whose strategy might also be a good fit for the available inventory. An auction then begins for a set period of time in which DSPs make higher proposals for the ad space. When the auction ends, the second-highest bidder will be awarded the space and be asked to pay around a cent more than the bidding price agreed. This model

is used so that a single Trading Desk cannot control the auction by buying all of the inventory that matches their strategy. If the highest bidder was allowed to win the auction, the biggest spending Advertisers would be able to pick off the best quality inventory.

The final agreed price pays the trading desk team, the use of the DSP, the access to the exchange and the Publisher fee for the ad space which would include cost of the use of the technology on the Publisher side such as the SSP and the Publisher ad server. On top of this, the Advertiser needs to also consider the cost of ad-serving.

Real Time Bidding (RTB)

The Publisher does not have to load volumes of available inventory via the SSP at a given time, but rather an available ad spot can be presented into the exchange **as soon as it becomes available**. This is completely transformative in digital advertising because it means that when a user arrives on a Publisher page, the available ad spot can load a tag belonging to the SSP, which channels the availability down to the exchange in real time. This opportunity is literally to present one ad, to one user, on one placement, on one Publisher in real time.

Essentially this is a yet un-fulfilled impression and as it loads in the exchange, the DSPs leap on it like fresh meat in a hungry piranha tank. The auction begins and ends in microseconds, the spot is paid for and the DSP connects the ad server ad tag to the SSP so that the ad arrives on the page as it loads. The user notices no difference between this method and any other method of ad delivery. This is referred to as Real Time Bidding or RTB.

To work most effectively the DSP strategy must already be set live and the SSP tags must be sitting on the dormant pages on the Publisher website. These SSP tags can be designed to carry extra information about both the user and the page down to the

exchange for real time bidding. If a user is logged in to the Publisher page, these details could be funnelled down along with page-level contextual information or the URL of the page. If the user information is funnelled down, it needs to have had the users consent before the ad space is offered out to a third party to adhere to privacy compliance. Such compliance however, may have been granted when the user signs up to get access to the Publisher site.

Cost reporting in the ad server

Apart from the delivery of ads from the ad server to the page, the DSP can facilitate the flow of reporting data from the exchange back into ad server reporting. This ensures that cost, delivery, verification and performance information is all still captured in a single system of record at the ad server. Since the ad server is counting based on a succinct definition of impressions, the connected cost information must be based on this methodology for the sake of data accuracy. This information can be channelled using data and API integrations between the ad server and the DSP. Advertisers and Media Agencies should connect with ad serving providers to understand more about this data flow.

DSP Strategies

When compared to buying ad space directly from Publishers and Networks, DSPs offer a greater access to available inventory ad strategies based on algorithmic optimization. A DSP will ensure it is not making bids on the same inventory twice. It will also make bids for inventory based on different cost models and calculate if the goals can be met with the available media budget, adjusting the strategy as it works its way through the budget so as to meet the goal.

Those bids which are lost to competitors contain insights which assist in winning bids for inventory with the same competitors in the future. This is referred to as "win and loss level data" and is a

proprietary data offering available to Advertisers from the DSP technology provider.

DSP Retargeting with DCO

As described in Chapter 7, the DSP strategy can incorporate retargeting information. For this to work the trading team can generate a retargeting tag within the DSP interface and send this to the trafficker working in the ad serving interface. The trafficker can piggyback the DSP retargeting tag onto a container tag or ad server retargeting tag sitting on the Advertiser site. When the user triggers the retargeting tags their ad server, the cookieID is added to the ad server retargeting pool but the DSP retargeting tag also drops a cookie on their browser. When the user arrives on a programmatic-enabled Publisher page, the SSP placement can be used to capture the cookieID for the DSP and pass it down into the consideration for the yet un-purchased impression in real time bidding. As the DSP recognises the cookieID, the corresponding strategy can activate a bid on the impression. If the DSP wins the auction, the ad server ad tag is delivered, first checking the ad server cookieID on the user's machine. Since there will be a record of this cookieID in the retargeting pool the ad server will serve the ad or version that corresponds to the user.




Figure 8.3 Dynamic retargeting with DCO purchased with programmatic. The user arrives on the Advertiser site, the versions in the ad server as built out on mass versioning from the Advertiser product XML. The same XML updates the Advertiser product page. The unique identifier for the product as an SKU number is passed into the retargeting tag for the ad server (R) as is associated with the cookieID on the ad server (2). The retargeting pixel for the DSP is piggybacked on the ad server retargeting pixel. This information at the cookie level for the DSP cookie is passed onto the DSP (3). The user then navigates to a Publisher page programmatically connected to an Advertising Exchange via an SSP (4). The placement loads into the ad exchange via the SSP (5) and (6). The DSP sees that the cookieID meets the retargeting strategy (7) and bids on the as yet unserved impression. Winning the auction, the DSP passes the ad tag onto the Publisher page. The ad tag reads the users ad server cookieID (8) and sees that it is matched to a dynamic variable as an SKU number. The ad server selects the corresponding ad version (9) and serves it into the Publisher placement (10).

Such strategies are employed by technology providers like *Criteo* and *Struq* who specialise in this specific kind of audience chasing. Unfortunately, with saturated audience chasing, such technologies will always win the conversion under the traditional last ad attribution model. Even if the user may have had every intention of buying the item without clicking such an ad, the convenience of the placement acts as a shortcut to Advertiser site. Marketers should think carefully about awarding full conversion attribution to such Publisher inventory.

DSP Keyword Targeting

An alternative strategy to retargeting would be that when the bid is won that the SSP passes a keyword value from the page, through to the ad server via the DSP. As with keyword targeting the ad server can use this value to display a corresponding ad or ad version.

Negative Retargeting

The DSP gives the option to offer greater audience reach by ensuring users are targeted that have never seen an ad or landed on an Advertiser site before. The strategy setup in the DSP can state that if the user is not in a retargeting pool, that they should be the target of the strategy. This is enabled via the presence of a specialist retargeting technology called an "exclusion pixel". Exclusion pixels can be generated from any of the DSPs and work using the opposite methodology to retargeting pixels.

DSP delivery controls

Popular features formerly reserved for use in Publisher ad servers are available in the DSP such as geo-targeting, strategy prioritisation and frequency capping. The settings get quite granular so that the DSP can effectively distribute the marketing budget.

Building Private Exchanges

Advertisers and Agencies that have more money to spend and like the idea of real time bidding but would prefer access to more premium Inventory, can work with the DSPs to setup restricted or "Private Exchanges". These are invitation-only marketplaces designed to give Advertisers with more money to spend, a "first look" at the available premium inventory before it is offered out on a public exchange. Alternatively a set amount of inventory can be procured at an agreed price between the Advertiser and the Publisher first, and the marketplace can act as management system to decide which ads to show to audiences based on decisions about the state of the inventory. If an Advertiser has multiple brands, there is no commitment to supply specific creative to the inventory until the very last moment before impression delivery. This makes for a more efficient use of first access to a private arsenal of inventory.



Figure 8.4 Private Exchanges sit alongside open exchanges as options for Advertisers and trading teams with enough budget to support the availability of the private marketplace.

Strategies based on Publisher Audience Descriptors

Since Publishers, Networks and SSPs bring the inventory to the marketplace and label it when it arrives; the Advertiser-side trading teams are not making decisions about the media space based on an independent system of audience classification. A Publisher may well have dressed up the inventory to make it more attractive to sell and the Advertiser can end up displaying an ad on a page which has little or no match to the user seeing the ad. Trading experience and historical reporting on ad performance

will demonstrate which of these Publishers are classifying their audience ineffectively. But the demand for a good match every time has seen the emergence of systems that bring more information to the trade. The collection and presentation of this information can be facilitated by a third party, which can establish more value to the inventory in the marketplace.

Cookie Level Data Providers

Trading teams can buy audience data from specialist third party data providers such as *VisualDNA*, and *BlueKai*. Alternatively data is available from collectors of traditional digital data points such as email addresses from providers such as *Axiom*. These providers have datasets about audience types based on cookie level information. The Advertiser can use this dataset to look up more information about the user of an un-purchased impression before the decision to bid on the impression is made.

In addition to using third party data, the Advertiser can use their own datasets from systems like their CRM to get a better understanding about the un-purchased impression. This is referred to as first party data and is considered by many to be the most valuable data that an Advertiser can collect and activate for use in the programmatic space. Providing there is a cookieID attached to the dataset, the data can be used. To do this, the dataset is synchronised with the available cookies on a user's machine in a complex process called "cookie matching" and is facilitated by an intelligent data warehouse system called a DMP or Data Management Platform.

Data Management Platforms (DMP)

DMPs are easily plugged into the DSPs to inform the bidding strategy and some DSPs come with their own DMP facilitates. Players include platforms such as *Audience Science* and *eXelate*. These platforms are growing in popularity as they provide more effective decision making on the buying of inventory in

programmatic by improving the ROI. Data activated via a DMP might show that a user is on the cusp of making a purchase. With the DMP activated, programmatic facilitates the opportunity for an Advertiser to pay a relatively high price (once data costs are factored in) to get the ad in front of the user. If the purchase that this triggers offsets costs then the DMP data enablement strategy becomes a particularly lucrative one and the data can act as a competitive advantage over other bidding Advertisers in the marketplace. Examples include providers that specialise in optimization such as *RocketFuel*.




Figure 8.5 DMPs can offer value in multiple positions in the buying and planning cycle within programmatic. Data that can be matched to existing cookieIDs at the user level add more information about the user to the Advertiser buying decision and will allow the trading team to adjust the buying strategy accordingly. Without additional data plugged in via the DMP, the Advertiser is almost buying blind, taking the existing audience classification as truth without independent verification.

DMPs are data storage and activation facilities for data that Publishers can also utilise. In buying or hiring third party data, the Publisher can independently verify the quality of their inventory as it is loaded into the marketplace to make it more attractive to Advertisers. The cost for doing so can be set in the minimum bid price that the Publisher is willing to set as the auction begins.

Page Level Data Providers

Outside of the cookie space, data can also be plugged into the DSP that does not depend on user based information but instead on page level information. Examples include *GroupM's Crystal Semantics* or *Peer39*.

As explained in Chapter 5 - these technologies can be used to pre-emptively verify that the page in question is a suitable environment in which to display an ad. The DSP can incorporate this information into the buying strategy to target towards or away from the defined variables. DSPs will also typically enable traders to load blacklists of sites as part of their strategies to ensure that banned sites or pages are never considered within the marketplace.

Regulatory bodies encourage DSPs, exchanges and companies in programmatic to facilitate the safe purchase of inventory such that universally recognised banned sites and blacklists are adhered to as part of a crackdown against fraud and cybercrime.

Channel-Specific Complications in RTB

In the current landscape, inventory can be bought for almost every media channel through the exchanges. There are specific issues related to RTB for Video and Mobile. Mobile suffers because an impression loaded into an exchange in real time may not necessarily be able to deliver the ad before the mobile signal drops out. Video inventory suffers from a problem of supply where there is high demand for Instream video ads but limited sources of quality content. This is why ad server verification and fraud reporting is in high demand within programmatic. The allure of high pricing on Instream inventory has led to some long tail Publishers engaging in fraudulent activity to bury Advertiser inventory among poor content or into spots that were misrepresented.

DSP - ad server of the Future?

DSPs and ad server have not merged as technology providers due to the impartiality offered by the ad server in its total separation from the interests of the Publisher. Since the ad server independently declares how many ads were tracked or delivered, the ad server reporting shows the Advertiser how much to pay the Publishers and it is of paramount importance to keep these technologies separated.

Third party ad servers can continue to offer de-duplication at the cookie level and provide rich channel ad serving for DCO while encumbering Search, Mobile, Instream, Affiliate and Display. The reporting alone in pathway to conversion, cookie level data, adjustments for cookie deletion, post-delivery verification and viewable impressions means additional first party data for the Advertiser to activate inside the DSP. However, the third party ad server remains the hub of creative optimization where the DSP is emerging as the hub of the media optimization in the future of digital advertising.

Chapter 9

The Big Picture


Chapter 9

The Big Picture

- *Figure 9.1 - Display and Programmatic*
- *Figure 9.2 - Layering in Paid Search*
- *Figure 9.3 - Layering in Affiliates*
- *Figure 9.4 - Layering in Users, Advertiser Properties, Ads and Trafficking*
- *Figure 9.5 - Layering in Retargeting*
- *Figure 9.6 - Third Party Ad Serving - the Diagram*

Having covered the aspects of ad serving at a high level, it is possible to piece the management of the channels together to get a third party ad server perspective of digital advertising. The first diagram picks up from Chapter 8 by incorporating Display and Programmatic.


Figure 9.1 - Display and Programmatic



The Media Agency or Advertiser can procure media inventory either through the direct buy channel; speaking to Publishers and networks directly to make a deal on the ad space or rules can be set up in the DSP along with a budget to bid on inventory in the exchange environment. The DMPs add context to the impressions and ad spaces being purchased with the option to trade in real time. This same method can be applied to all Display channels

(Standard Display, Rich Media, DCO, Instream video and mobile). The third party ad server plays a role in supplying the creative as an ad tag to the DSP or direct to Publisher to be implemented into the Publisher ad server.

Figure 9.2 - Layering in Paid Search




Search has matured programmatically faster than Display. The established technologies sit with the Search engines as SEM systems like *Adwords* and as bid management solutions. The third party ad server offers tracking integration to each to have eyes on these channels, while offering cross channel data through a pathway to conversion feed to add intelligence to the bids on keywords.

Figure 9.3 - Layering in Affiliates




Since affiliates also utilise search there is a line in to the same SEM tools. The third party ad server also offers click based tracking to the affiliate. Affiliate reporting is fed back into the ad server.

Figure 9.4 - Layering in Users, Advertiser Properties, Ads and Trafficking



Introducing basic views on the user perspective and Advertiser perspective; The user visits Publisher pages, properties, Search Engines or affiliates to utilise content via the browser and device, carrying a publicised IP address and, providing they have not opted out of third party advertising, a third party cookie. The user is delivered an ad via a placement on the Publisher site. The ad server counts the delivery and the creative is displayed. Should the user interact or click, the ad server counts the actions and redirects the user via the clickthrough URL. The user arrives at the destination page such as a landing page. The landing page and subsequent pages can carry retargeting tags and conversion or container tags. The conversion tag feeds the user conversion event into ad server reporting. The analyst can read the reports and provide insights based on behalf of the Media Agency or the Advertiser

Figure 9.5 - Layering in Retargeting

The retargeting tag can be used to segment an audience or cookie pool. The data can be used for creative optimization in the ad server or pushed to the Demand Side Platform as a data set to target audiences for media optimization.

Closing Chapter

Third party ad serving technology is a marketing revelation.

Far from being a tracking and serving technology, the data that the third party ad server produces is an industry backed independent record for Advertisers to establish a security around the purchase of digital media space.

It is relevant to draw attention to the apparent parallel which exists between banking and ad serving. The attitude is one of trust. Housing the data which tells the Advertisers what to pay the pay Publishers; this highlights the value of account-keeping across a single methodology with global scaling, a level of commercialisation required to support the global communications system we call the Internet.

Taking a step back from the technicalities and evolving features, the foundation laid by the third party ad server brings economical wealth to the digital stakeholder groups: Creative Agencies, Media Agencies, Advertisers and Publishers. And as this technology diversifies and changes, it becomes ever more important that as an industry, we remain educated and informed; challenging the foundations of Advertising so that it's benefits always nourish the end user above all else. Advertising as a whole is not always embraced by consumers but its value in providing a free internet for all must mean that users and Advertisers can reach a mutual partnership.

This book is designed to make the technology side of digital advertising more accessible to everyone. The concepts covered here may seem fresh and new but digital evolves by the day. To extend the value of advertising technology further involves

reaching forward towards the goal of better advertising effectiveness. And Advertising rightfully must lay its claim in a meaningful and scientific way on the outcome of a purchase as a direct influence over the behaviour of the consumer.

Business efficiency demands that the general art of indirectly influencing a purchase be made fully accountable to prove that general advertising is not a waste of money. As we have seen, ad serving technologies and new additions to the scene like DSPs and other point solutions are gradually filling in the data gaps. However if advertising were truly effective there would be no doubt as to its capabilities and all profits would be diverted to the cause, to generate ever more sales until brands and their products and services saturate the consumer space.

Saturation is a genuine problem, far reaching campaigns that overlap one another have a maze of data to process to realise effectiveness and cement true attribution. The race to track all channels reaches out into the TV and video space as mediums converge and the user consumption of advertising evolves in new ways. And where TV and Online teams both meet, the goal remains the same. The technology, it is perceived, if used correctly, can do what advertising has attempted to do since its advent: correctly predict and drive user action to purchase indefinitely.

Today's third party ad servers are gathering the data that marketers need to realise this dream but the information remains in silos. In the economic competition of advertising, the data sits with the Advertiser; their successes are their secrets and the best stories told in the best performing data, feed cleverly into the sales algorithms of marketers. Ad serving technology is the enabler to this end and is the metal detector revealing the hidden treasure of this success.

For it is the long labour, cost and learning of the marketer, that bring the fruits of such efforts to the possibility of such a future. But we are only at the start of the journey. Hopefully this book invites greater opportunity to those eager to shape the bright future that lies ahead for digital advertising.

Thank you for taking the time to read this book. Once again I welcome any feedback from you, and I will endeavour to assist in any way I can to improving the industry for everyone through my work.

Further Reading and Study

This book has been carefully compiled of ad serving and advertising technology fundamentals. The detail contained within these pages could be classed at an 'intermediate' level of knowledge. Including the whole swath of upcoming trends, custom technologies and start-ups would have made this book too extensive to focus on those fundamentals. This is the framework from which to build new platforms, concepts and see how to find evolving opportunities in these technologies. As I was writing this book I was still learning new things and I have actively omitted some concepts which are more advanced. Truly I could have been writing forever and never published if I had included everything I know about Digital Advertising.

To help your learning beyond the fundamentals or to be trained on any of the concepts in this book in person please visit my website: <http://cristalconsultancygroup.com> here you will be able to attend training courses, receive private tuition and sign up for consultations for your business. Advertised training courses tend to be based out of London, UK but I do travel frequently so get in contact if you would like me to arrange training overseas.

..and to stay on top of the very cutting edge in Online Advertising Technology, beyond the fundamentals, please subscribe to my blog, published on the same website:
<http://cristalconsultancygroup.com>

Acknowledgements

This book took me a long time to write, I was not aware of the enormous amount of time I would have to dedicate to finally get it out 'into the wild'. Although there are very few references to outside academic material (mostly due to the infancy of the industry and technology) I want to very sincerely thank those that have helped educate me up to this level of knowledge. The list of names exceeds a hundred people, but at risk of missing off any names I just want to send a huge thank you.

In addition there have been whole rafters of proof readers, editors, minor contributors and supporters that have my unrelenting thanks. The efforts we undertake every day assures the free, public availability of the information technology we call the internet. For this there can be no greater reward than to keep our industry knowledge up to date and spend the days of our lives giving birth to inspiring new ideas, building new intelligent technologies and contributing directly to the evolution of our civilisation for generations of *mad men* and *women* to come. I wish you all eternal prosperity in heart and mind.

Finally I would like to thank my family, who have always believed in my capabilities and have supported my passion for writing and technology at every turn.

Ruby Hirsh, Ashley Cristal, James Cristal, Nick Cristal, Michelle Miller

Glossary

Term	Definition
302	An HTTP code to tell the browser that the page in question had been moved. This is called a server side redirect and is used by ad servers as a method to count and track ad activity.
third party code	Code belonging to an external system to the one being used. Third party code is typically used in a first party system. Such as example is where a Publisher uses a Publisher ad server and brings third party ad server tags into the system. The tags contain third party code.
404	An HTTP code to tell the browser that the requested page had not been found. 404 errors typically come up in ad serving when a click through URL lands the user on a page that doesn't exist.
4th party code	Code belongs to another technology provider from the third party provider. An example is where code is generated from a survey vendor and provided to the trafficker to implement on an ad such that when the ad closes the survey is launched from the fourth party survey vendor system.
Action tag	Another name for a conversion tag from the Atlas ad server. Action tag might be a more accurate description since conversion tags do not just measure

	conversions but can also measure non-conversion events: also called "actions".
Actions	A conversion tag loads whenever the code is rendered by a browser. If the user did not encounter an ad previous to the conversion the ad server does not register a conversion but makes a log of the code load. This is called an Action and an Action is not tied to a unique cookieID.
Ad assignment	The process undertaken by the trafficker inside the third party ad serving interface to assign ads to placements.
Ad brokering	A Publisher, network or exchange may be selling inventory with the claim that it falls under their ownership but in fact they are acting as a middleman to another Publisher network or exchange to procure the inventory at a certain cost. Be wary of ad brokering. More middlemen mean more technologies, greater discrepancies, larger margin for delivery error and hidden costs. Verification can provide some level of control but the industry is rife with the practice given the ease of passing ad tags onto an additional party.
Ad clutter	A Publisher that swamps web pages in ads clutters the page and reduces the value of the content and subsequently the ad space. Get a commitment from the Publisher that ads will be delivered to clutter-free pages and utilize advanced verification technology to detect for ineffective impressions under the guise of

	ad fraud.
Ad controls	The name given the settings on the placement which allow ads to be rotated, sequenced, targeted and delivered at a given frequency governed by a set of rules setup by the trafficker. If there is only one ad assigned to the placement then ad controls are limited.
Ad location	Once implemented on a Publisher site the ad appears to live in a fixed location on the web page when it loads into the browser. Although intrinsic to standard display ads, Rich Media ads can stretch away from the original ad location into an expanded state or not have a fixed location at all (such as with a Floating ad). The ad is essentially anchored in the ad location which from a Publisher's perspective is sometimes what is referred to as the "placement".
Ad server macro	Macros are small tokens or tiny pieces of code usually a single word long surrounded by symbols. The macro is an instruction to the ad server to load additional data into the position of the macro. This will typically be an identifier such as a Placement ID.
Ad serving Hierarchy	The tree-like folder system within an ad server in which campaigns are set up.
Ad spot	(see ad location)
Ad Tags	Once ads have been assigned to placements inside the ad server campaign

	management interface by the trafficker, the code can be sent to the Publishers. The code for each specific placement is called an Ad Tag.
Ad blocking	There are two types of ad blocking. One is performed locally on a client device usually within the browser to block an ad loading. The other is within a verification tool to prevent an ad delivery based on the ineffectiveness of the site or environment in which it is being requested to load.
Affiliate Auditing	Affiliates within a network or subscribed to a program are checked regularly by the team managing the program to ensure that the Affiliate site is still suitable to be included.
Affiliate program	An ad campaign extended across multiple Affiliate sites usually managed by an Affiliate network. The campaign creative and click-through for the affiliate program are available from the affiliate network or the program manager working at the Agency or Advertiser.
Affiliate Screening	A process undertaken by the Affiliate program manager or the Affiliate network to check that an Affiliate is suitable to join an Affiliate program.
Aggregated reports	Reporting templates constructed from ad server analytics and counting data, amalgamated across ad server entities so that reporting is a summary of activity and can be analyzed and understood easily

	without the specialist skill of an analyst.
API	Application Programming Interface
ATD	Agency Trading Desk
Behavioral Targeting	Another name given to Retargeting.
Black Box Technology	Not disclosing the workings of a technology owing to the complexity or implication of doing so. An example would be revealing Google's search algorithm. Google as a search engine is a black box technology. The workings remain a mystery but it is agreed that it does the job. Revealing the workings would allow others to defraud the system.
Blind Networks	To reduce the cost of inventory, planners and buyers can purchase inventory based purely on audience information and not on the name of the site where the final ad will be delivered. This gives Publishers and Networks more flexibility to optimize delivery of ads across Networks based on inventory availability, rather than having to commit to the final location earlier on. To ensure the audience matches the Advertiser requirement, verification reporting and ad blocking can ensure the page is brand safe on ad delivery.
Blind Networks	A blind network is Publisher network where inventory is labeled only by the audience type such that the final location of the delivered ad is not known to the Advertiser. Blind Networks typically sell

	inventory at a lower cost so as not to have to commit to the exact page location of the ad before it is known if that spot will be available.
Booking pathway (on an Advertiser site)	The design of the site to drive users from the point of arrival (the landing page) to the point of conversion (the thank you page).
Brand Advertisers	Advertisers whose main goal is deliver a creative message. This might be to create awareness rather than to direct traffic for the precise goal of making a sale. The sale event may happen a long time after the brand activity of the ad campaign. This Advertiser is focused on delivery goals and GRP. Common Brand Advertiser verticals are FMCG and Autos.
Cache	A machine will store copies of files and information to make the retrieval of that information faster if it is requested again. A cache should be refreshed often so that the user does not get out of date information. This applies to ads and websites and a cache can occur locally on a user's device or in the cloud.
Call-to-action	A part of an ad or web page designed to be engaged with by the user. This could be a submit button or trigger.
CDN	Content Delivery Network
Channel	A channel within the context of digital marketing consists of any sub-channel within Digital. This also spreads itself to

	format complexity such that the most popularly referenced channels are: Standard Display, Rich Media, Search (SEM), Natural Search (SEO), Dynamic Creative (DCO), Mobile, Instream Video, Affiliate, Social and Email.
CLD	Cookie Level Data
Clean Data	A consistent naming methodology produces large sets of data that can easily be sorted, concatenated or normalised. This is clean data and is used in reference to utilises consistent naming conventions as ad, campaign or conversion tag names.
CMS	Content Management System
Code Wrapper	Code which sits either side of the main subject to effectively wrap the subject up.
Companion Ads	An ad that sits within visible range of an Instream video ad to accompany it.
Container Tags	Conversion tags which are capable of having other third or 4th party tags or pixels piggybacked on top. The conversion tag therefore acts as container to the other tags. In the DoubleClick ad server these are referred to as Floodlight tags.
Contextual targeting	Also called Semantic targeting. The placement code extracts phrases and words on a page to understand the subject matter of the page in question before deciding to deliver an ad based on the interpretation of that information.
Contracted State	An expandable ad that is dormant or has

	already completed its expansion effect.
Conversion Event	A conversion event is just another way of describing the event that is the conversion. Conversion event is a more universal term than "conversion" which has a connotation that the event in question was a completed sale.
Conversion tag	Generated within the ad serving interface, the conversion tag is a piece of code that can be placed in an ad or on Advertiser site. If a user has previously been encountered by the ad server they will trigger a conversion event in ad server reporting.
Conversions	A record of an event performed by the user, such as the loading of a specific page or ad which signifies the satisfactory completion of a marketing goal. This is most commonly a sale.
Cookie Matching	If two separate datasets both use cookies as the unique identifier for the user, the cookies can be matched together to see where a user is common to both databases. For instance an Advertiser could produce a script to site on the Advertiser page that can read the users first party cookie ID for the website. The same script could take the cookieID and pass it into the extended data of an ad server conversion tag. The cookie level data would reveal a match between the two cookieID's.
CPA	Cost Per Acquisition. A common model to

	pay Publishers based on their effectiveness at driving users to convert or purchase based on the subject of the ad.
CPC	Cost Per Click. A common model to pay Publishers based on the user clicking on the ad. CPC is most effectively used in the Paid Search channel.
CPL	Cost per Lead. Another name for CPA. The conversion event in question is the collection of user data.
CPM	Cost per "mille". Cost per thousand impressions measured. The most common cost model in digital advertising.
CPQ	Cost per quote. Another name for CPA.
Crawlers (Search engine Crawlers)	See "Spiders"
Creative Agency	An external provider of design and development services to build code and deliver the creative which eventually becomes the ad.
Creative Fatigue	Most commonly a failure to implement frequency capping so that the creative message is shown so often that the user sighs in disappointment upon seeing the ad again and again with no control to stop its repeat appearance around the web.
Creative shop	An umbrella term to mean Creative Workshop and will mean a Creative Agency or an internal creative team at a Media Agency. The Creative Shop is home to the creative designer and developer

	roles.
Creative team	A reference to the creative developer and designer roles that can sit within a Media Agency or at a Creative Agency. The creative team is responsible for the design, development, build and delivery of the ad creative and may also be involved in the design and maintenance of the Advertiser website.
CRM	Client Relationship Management system. An example would be <i>Salesforce</i> or <i>Microsoft Dynamics</i> .
CTR	Click-Through Rate. A percentage of the number of user that saw (incremented an impression) and then clicked on the ad.
CV Tools	Content Verification Tools.
Data Leakage	Publisher audiences can be harvested using retargeting tags. The harvesting of these details for further use later on is considered a leak of that Publishers audience. The claim may be that the audience is the property of the Publisher.
DCO	Dynamic Creative Optimization.
De-duplication	Always making sure that the entire audience for all ads belonging to an Advertiser is tracked in a single system such that users can be recognized as unique and their unique pathways to events such as conversions can be described in the ad server without any question that some data belongs to the

	same user more than once.
Delivery	A reference to ads being delivered by the ad server. A delivery is measured in impressions so delivery relates to impression based goals and data. Delivery of ads online can be compared to TV ads being delivered to TV viewers. The ad is simply sent to the end recipient with no confirmation that it was ever seen.
Desktop Device	A desktop device is a name given to a computer that remains in a fixed position such as a PC. A tablet is not classed as a desktop device but a laptop may fall into this category.
Device Fingerprinting	For mobile advertising, devices do not carry a universal unique identifier such as a third party cookie in every case. A device fingerprint is an identifier used in place of a cookie for the purposes of attribution and the measure of unique reach and frequency.
Direct Response Advertisers	Advertisers more focused on driving conversion events or sales. This Advertiser is focused on performance channels like Search and Affiliate as well as performance goals like clicks and ROI. Common direct response verticals are online retail and travel.
Display Ads	Also called Standard Display ads. These are the most predominant visual ads as flash seen across the web.
Domain	The main folder in which a website

	resides. The domain sits at the top level and subdomains sit within it.
DPB	Data pass back - the <i>Atlas</i> ad server version of Cookie Level Data.
Drop off rate:	The diminishing percentage of users arriving to pages. Those that arrive on the landing page but fail to convert on the thank you page are said to have dropped out of the booking pathway.
DSP	Demand Side Platform
E2C	Exposure to conversion reporting - the DoubleClick equivalent of P2C. See P2C reporting.
E-conversion	Dividing the attributed conversion among touch-points on the pathway to conversion allows each touch-point to be valued. This assigned value is rarely a whole number and is referred in the <i>Atlas</i> ad server as an e-conversion.
Effective Impressions	Impressions that met the Advertiser criteria for quality verification.
Entities	Ad server entities are named parts of the ad server hierarchy that will typically have a corresponding entity ID such as Advertiser, Placement, Site, Channel, Ad, Creative and Version.
Events in the pathway (to conversion)	Also called "touch points". User see, click and interact with ads for an Advertiser before finally converting. These ad engagements are referred to as non-conversion events on the pathway to

	conversion.
Exclusion Pixel	A retargeting pixel specifically used to negatively retarget against the cookie pool it collects. Predominantly used by the DSP.
Execution	The Creative is sometimes referred to as a "creative execution" or just an "execution". Another way of saying the final delivered ad.
Expanded state	A Rich media expandable ad, at its most full dimensions maintained for a period of time or until a user engages to reduce the size, changing the state of the ad to the contracted state.
Experiment	(for user targeting testing like AB testing);
Factors Report	A report that can recalculate attributed conversions based on different attribution models from the one used to calculate the conversions in the first instance.
File weight	The size of the file. For ads this is measured in kilobytes (kb) and megabytes (mb).
Flash Cookies	A storage facility within Flash to remember information about the user and their relationship with the ad for the information to be used later on. Flash cookies are not favored in Digital advertising because flash cookies are hidden from the user as a default and are more difficult to remove from a user's device.

Floodlight tags	<i>DoubleClick's</i> name for a container tag. See container tags.
FLV	Flash Video. A reference to the file type that creative developers supply to play video within an ad.
Frequency	The number of times that a unique user is exposed to an ad. Can be capped by frequency caps.
FTP	File transfer protocol. Referenced with regards to large data sets that can be posted to an Advertiser's FTP site for the purposes of large file delivery and storage from an external system such as the third party ad server reporting system.
GDN	The <i>Google Display Network</i>
Geo Performance	Also called Geo Reporting. The name given to a report that shows user IP address mapped to a location from country level down to postal and zip codes in some regions with varying accuracy. Geo performance is used to ensure Publishers are correctly targeting ads to users in the agreed geography.
GRP	Gross Rating Point. A measure from TV advertising to extrapolate effective reach and frequency. A GRP number can be so refined and calculated so as to accurately point at a given number of in-store sales if the goal is reached.
Hierarchy	The tree like folder structure housing the ads in the ad server interface. Advertisers,

	campaigns/media plans, networks, Publishers, sites, sections, packages,
HTML5	The latest edition of the mark-up language HTML with improvements to support video and JavaScript.
IAB	Internet Advertising Bureau
Impulse conversions	A user conversion which occurs within a single ad exposure and within a short time frame.
In-App	Ads that appear within mobile applications.
Ineffective impressions	Impressions that failed to meet the Advertiser criteria for quality verification.
Inventory	The amount of advertising space available.
ISP	Internet Service Provider
ITD	Independent Trading Desk
Jump Tag	A DFA click command that can be used to carry the user to an over-ride page URL.
KPI	Key Point Indicator. A marketing term used to describe a goal set by the marketing team seen as a priority for success.
Media Buy	The purchased ad space.
Media Plan (ad-server hierarchy entity)	Within the third party ad serving interface, the media plan is the name given the campaign level.
Media Plan	The documentation that lists all of the

(Agency document)	media buys.
Negative Retargeting	See Exclusion Pixel.
Optimization	Using the reporting data to make decisions to improve campaign performance, delivery and verification.
Order ID	The unique identifier for the conversion event generated by the Advertiser and passed into a conversion tag on the Advertiser page.
P2C reporting.	A pathway to conversion report. Typically a cookie level data feed showing how each cookieID encounters events or touch-points as ads before reaching a conversion tag with timestamps. A P2C report allows the Advertiser to construct attribution models. P2C is the acronym given to the <i>DG/Mediamind</i> cookie level pathway to conversion report.
Parallax	A creative effect common to eastern cartoon where background and foreground move at different speeds. The effect requires a high capacity of user RAM and so is only being adopted recently.
Pathway to Conversion	See "P2C Reporting".
Payload	The payload is the content of a token which could be an ad server macro. For example: "?rtu=payload".

PCP	A Publisher custom parameter. A token which can be populated in the ad call or click on the Publisher side to populate ad server reporting with more information such as a player ID.
Performance	A reference to non-delivery data about the relationship between ads and users received and processed by the ad server. Performance is measured in clicks, conversions and engagements/interactions so performance relates to goals and data based on these metrics. Performance is a feedback mechanism, allowing the analyst to determine if ads effectively met goals related clicks, conversions or engagement.
PII	Personally Identifiable Information
Placement	Just like with the word "tag", the word "placement" has dual meanings. From the perspective of an Advertiser and third party ad server, placement means the code package that is sent to the Publisher and from which the ads or creative called to serve from within it, are played. Placement in this context is the portable housing in which the ads in rotation or sequence reside. From a Publisher perspective "placement" refers to the fixed ad spot on a page which never moves. Traffickers supply to the Publisher ad operations team code as an "ad tag" which when inputted into the Publisher ad server are eventually served into the Publisher's 'placements' all around the

	Publisher site.
Point solutions	Specialist technology solutions typically with advanced features in the field of their expertise. Utilizing a point solution in addition to an ad server will incur higher costs and integrations between point solutions and the ad server are paramount to ensuring clean data. Many such technologies are start-ups. Question the ability to scale to multiple markets, and the extent of the available support.
POS	Point of Sale system.
Post Click Conversion	A conversion where the attribution model in place awards this conversion to the user clicking.
Post Impression Conversion	A conversion where the attribution model in place awards this conversion to the user being exposed to an ad.
PPC	Pay Per Click. Another name for CPC
Pre-emptive Verification	Technology which checks the yet-un-purchased user impression at the page level to determine if the environment is suitable for the Advertiser using the DSP to proceed with the auction.
Premium Inventory	Ad space which is in high demand and comes at a high cost relative to remnant inventory.
Publisher Audiences	Publishers will sell their ad space based on the marketers request for a matching audience. The ad space is labelled up and sold by describing the audience that will

	be exposed to the ads that will eventually occupy the space.
Publisher Property	A Publisher website is part of the Publisher's property as is a Publisher's app. This phrase is most commonly used as an alternative to describing a single part of the Publishers
Publisher Specs	Publisher requirements and limitations for the development of ads that will be permitted to be delivered on their site or properties.
QA	Quality Assurance. The period of testing and ad to check that it fits with ad server or Publisher specifications.
Query String	A token appended to ad tag on either the click or impression portion to carry out an additional function. Query relates to a request for information from a database.
R&F	A shortened form in reference to Reach and Frequency
Reach	The total number of unique impressions.
Recency	This describes the order by which things happened with the most recent in the primary position. The term is mostly used when describing custom attribution modeling in order to place a value on events that occurred more recently.
Recordability rate	See VTR.
Referral URL	The URL of the page before the page the user is on. The referral URL is counted as

	the URL responsible for driving the user to the present URL. The referral URL can be the page the ad was shown on.
Re-messaging	Another name given to Retargeting.
Remnant Inventory	Ad space which is not popular with buyers and is available at high volume for low cost from a Publisher.
Rich media specs	A common term used by creative agencies and developers to ask what coding requirements and restrictions the ad server will enforce to ensure ads pass QA.
ROI	Return on Investment. A marketing/business term to establish what the Advertiser gets in return for the marketing spend. Intelligent marketers can use ad serving technology to establish ROI goals and see if advertising efforts are generating incremental sales or reach.
RTB	Real time bidding. In reference to programmatic buying.
SDK	Software Development Kit. In reference to mobile application development.
Second Screen	The user may be using one device to consume media at the same time as another. The primary screen attracts most of the user attention, the second has less. User trends suggest users consume social media on the second screen while consuming premium video content on the primary screen.
See through rate	The percentage of user impressions that

(STR)	could be checked for Verification.
SEM	Search Engine Marketing or Paid Search. Also used interchangeably to describe the keyword creation and buying toolsets (such as <i>Google Adwords</i>) or used to describe a bid management tool (such as <i>Marin</i>)
SEO	Search Engine Optimization or Natural Search. A non-paid for advertising channel where a website is optimized by the site developer to attract more search engine user traffic.
Single-counting Vs. Double-counting	A setting at an account level which allows conversions to be de-duplicated (single counting) or duplicated (double counting). Predominantly used to control attribution reporting for Natural Search.
SKU number	Also called a Product ID. Used in reference to DCO.
Spiders (Search Engine Spiders)	Spiders are programs that crawl the web, simulating a user, loading pages, digesting the information those pages and passing the information back to the source database for cataloguing. In Search Engine Optimization these technologies are called Spiders.
Spotlight tag	A term from the <i>DoubleClick</i> ad server to mean a conversion tag that does not have container functionality.
SSP	Supply Side Platform. In reference to programmatic.

Staging environment	A test page or series of pages on the Publisher or Advertiser side which tends to receive traffic by invitation to test ads and content before putting the live on the main website. The staging environment is controlled by the CMS and is a term common to tag management and ad QA.
Subdomain	The prefix before the domain or directly after it. The subdomain is designed to depict an area of a site or domain that might indicate a site section. Subdomain is referenced in ad serving as part of domainURL verification, reporting subdomains where page level transparency is unavailable to call out ineffective impressions that may impede verification such as brand safety.
SWF	The file type of flash animations. Most non-mobile display ads are SWF files.
TMS	Tag Management System or Universal Tagging System.
The Cross Domain File	A security feature of websites such that the implementation of the file allows for the reading and writing to remote domains. In reference to DCO.
The Fold	The fold of the user's browser is the horizontal or vertical cutoff where the remainder of the ad or web page is not visible. Scrolling the page reveals content or ads previously hidden by the fold. A term commonly used to describe the functionality of Viewability.

Token	An addition to an ad tag to perform a function such as carry additional information to another system. See "Payload".
Trafficker	The role which manages the campaign setup inside the third party ad serving interface. Traffickers are usually employed at a Media Agency or on-site at the ad serving provider. Traffickers can also be hired to work remotely via a trafficking outsource company.
Universal Tagging	See Tag Management System
URL	A web address for HTTP
User Agent	The technical term for a user's browser on any device. The user agent is a code passed to the ad server to identify the browser type.
Version	A reference to a variant of an ad for DCO. Each version is essentially an ad in its own right but would have been generated by a machine rather than created by a creative developer.
View through rate (VTR)	The percentage of impressions for which Viewability could be measured.
VOD	An overused incorrect industry reference to the Instream Video channel.
Wastage	A reference to impressions delivered ineffectively such that the delivery of the ads will be irrelevant to the user and could not possibly be used to attribute the ad to a marketing goal. The marketing

	budget is wasted.
White Labeling	Disguising technology ownership such that a technology appears to be owned by one party but in fact is licensed from another vendor.
XML file	In reference to DCO. An additional file to accompany an ad to allow the ad to be updated without requiring a flash developer to do so.

FIN

Made in the USA
Las Vegas, NV
29 May 2023



72655335R20331

Ad Serving provides a commercial infrastructure to the internet, spanning all sites, and touching all users, all the time.

Ad Serving gives marketers the ability to deliver and measure their ads for exceptionally low cost, revealing the true value of the ad space they pay for.

Utilizing cookie and page-level analysis, marketers can build vast pools of pseudonymous data about websites, ads and the users that encounter them. This data, and the trusted publisher-independent methodology for measurement, secures investment for Publishers. The counting, tracking and the delivery of ads to enable this data takes place at massive scale. The processes and systems behind data collection are complex, and marketers are expected to learn how to use them as the industry grows; applying this knowledge to their new roles in Digital Advertising.

Many undertake this challenge, unaware of the complexity that lies ahead. The learning resources are few and far between. This book satisfies an intermediate-level of tuition to Ad Serving Technology, illustrating how and why Advertisers continue to grow their ad spend in the planet's favorite new media channel.

ISBN 9781484867570



9 781484 867570

90000 >

