Ad Serving Technology

Understand the marketing revelation that commercialized the internet



POSTSCRIPT 1012

Gregory Cristal

Ad Serving Technology

Understand the marketing revelation that commercialized the Internet

Gregory Cristal

Copyright © 2014 Gregory Cristal

Cover Design by Andreas Brooks

All rights reserved.

ISBN: 1484867572
ISBN-13: 978-1484867570

CONTENTS

	Section Name	Page
	Purpose - Why write a book about third party ad serving?	25
	Part A - Setting the Stage	27
	The flow of this book	28
	What is third party ad serving?	30
	Getting Some Perspective	31
	Users	31
	Publisher	31
	Agencies	32
	The Advertiser	32
	Definitions and Methodology	33
1	**Chapter 1 - Section 1 - Browsers and Devices**	35
	The role of the browser in Digital advertising	37
	Figure 1.1	38
	Figure 1.2	40
	Tracking across multiple browsers	40
	first party Vs. third Party Cookies	41
	Figure 1.3	42
	Cookie Deletion	43
	Distinction between Devices and Computers	44
	Ad blocking Tools	46
	Why is there advertising on the Internet?	46
	Chapter 1 - Section 2 - The Anatomy of Sites and Ads	49
	What the user sees	49
	Figure 1.4	50
	HTML	50
	Figure 1.5	51
	JavaScript	51

Images 52
Text Links 53
Figure 1.6 54
Flash 54
The advent of HTML5 55

Chapter 1 - Section 3 - Ad Serving 57
What does the ad server do? 57
Figure 1.7 58
Table of Definitions (Table) 58
Figure 1.8 61
The role of the CDN 61
Waiting around for Ad Delivery - the user experience 61
Pre-ad display 62
During ad display 62
On ad close 62

2 **Chapter 2 - The Industry Players** 65
The Agency holding groups 66
Who are the big Agencies? 67
Running a Media Agency 67
Creative Agencies 68
Figure 2.1 68
Publishers 69

Chapter 2 - Section 1 - Advertisers and Marketers 71
Who is the Advertiser? 71
The Marketer 72
Agency pitch 72
Relationship with the Media Agency 73
Relationship with the Creative Agency 73
Figure 2.2.1 75
Controlling the pages of the Advertiser Site - the CMS 75

Landing Page 76
Basket Page 77
Confirmation Page 77
Figure 2.2.2 78
CRM Platform 79
Onsite Behavioral data 79
Tracking the users on site - Site Side Analytics 79
Click to Land 79
Passing the tokens in the click-through URL 80
Referral 81

Chapter 2 - Section 2 - Agency Planning and Buying 83
The role of the Planner 84
How to plan? 84
Figure 2.3 85
Research Tools 85
Figure 2.4 86
Demographic Digital Media Consumption by Publisher Website 86
Unique Audience - Site Overlap 87
Research through Benchmarking 87
In-Agency planning solutions 88
Figure 2.5 90
Technology for Buying 90
Agency Trading Desks 91
Figure 2.6 92
The Agency-side Media Plan 92
Selecting an Ad Server 94
What technology-based Services do Media Agencies offer? 95
Media Agency data analysts 95

Chapter 2- Section 3 - Creative Shop 97

The Creative Shop 98
Relationship with the Advertiser 98
Relationship with Publishers 99
Relationship with the Media Agency 100
Marketing budget for Creative development 102
Roles and Responsibilities 102
Figure 2.7 105
Getting the Design right 105
Development and Production 106
Flash authoring 107
Figure 2.7.1 108
Actionscript 108
Figure 2.7.2 110
Components, building and publishing ads 111
Ad storage systems (Asset Management Systems) 112
Localization 113
Standard QA 114
Figure 2.7.3 115
Uploading Creative 115
What's in a name? 116
Creative Inspiration 117

Chapter 2 - Section 4 - Publishers 119
What are Publishers? 119
Figure 2.8.1 120
How do Publishers sell their ad space? 121
CPM 121
CPC 122
CPA 122
Publisher Networks 123
Figure 2.8.2 124
Niche Networks 125
Blind Networks 125

Publisher ad servers 126
Figure 2.8.3 127
Inventory 127
Publisher Specifications 128
Publisher QA 128
Publisher Ad Ops 129
Figure 2.8.4 130
Ad prioritization 130
Frequency capping 130
Over and Under Delivery 131
Browser Targeting 132
Figure 2.8.5 133
How does Publisher reporting differ from third party reports? 134
Discrepancies 134

Section B - Campaign Setup by Channel 137
Figure 2.8.6 139
How does the end to end Campaign Management workflow work? 139
Billing, Electronic IO and Ad serving rates 140

3 **Chapter 3 - Ad serving Operations** 143
Trafficking 144
The third party ad server user Interface 145
User Access for campaign setup/ Permissions 145
Using multiple third party ad-servers (double-tracking) 146
Figure 3.0.1 147
Figure 3.0.2 148
The golden rule of ad serving 149
How does the trafficker check the ads are working? (Setup verification using an HTTP tracer) 149

Chapter 3 - Section 2 - Display ad serving operations 151
Day in the life of the trafficker 153
Foundation Level trafficking for standard display 154
1. Creating a place in the ad server to setup the campaign 155
The third party ad server hierarchy 155
Figure 3.1 157
Settings at Agency and Advertiser levels - Silo Capabilities 157
Figure 3.2 159
Settings at Agency and Advertiser levels - Sharing Capabilities 159
2. Structuring the campaign setup to replicate the Publishers on the media plan 160
Media Plan 160
Figure 3.2.1 161
Trafficking Sheets for outsource trafficking 162
Sites 162
3. Identify the formats of the ads and creating the placements 164
Ad Formats 165
Banner 166
Leaderboard 167
Skyscraper 167
MPU 167
Roadblock and Skins 167
Placements and Packages 168
Cost Packages 168
4.Assign the creative and setting rotation controls 170
Naming Conventions 170
Ad & Creative 172
Automated QA 172
Ad Preview 173

Adding the click-through URL: clickTag 174
302 Redirects 175
Adding multiple click-through URLs 175
Ad classifications / Custom Categories and searching 176
Ad Duplication 178
Figure 3.2.2 179
Figure 3.2.3 180
Readying the Ads for Assignment 180
Backup gif 181
Possible Placement Assignments 182
Figure 3.2.4 182
Ad assignment 183
Multiple ads to a single placement 183
A single ad to multiple placements 183
A single ad to a single placement 184
Figure 3.2.5 185
Adjusting the Ad Controls 185
Frequency Capping 185
Ad Rotation 186
Sequencing 186
Weighting 186
Ad Prioritization 187
Time Based 187
Post Cap Gif 187
The two levels of control 188
Figure 3.2.6 189
5. Trafficking the Tags to the Publishers 190
Ad Tags and Tag Formats 190
Figure 3.2.7 191
Flighting 191
Thoroughly Checking an Ad is live 192
6. How does a trafficker manage campaigns with hundreds of Ads? 194

Bulk trafficking through *Excel* templates 194
Bulk trafficking with *Excel* as the UI 194
Automated trafficking setup via ad server API 195
Figure 3.2.7.1 196
The JavaScript carriage capability 197
Survey vendor Code and third party scripts-in ads 198
Figure 3.2.8 200
Tracking but not serving 201
Is it possible that a Publisher cannot support third party ad serving? 202
Figure 3.2.9 203
A note on Click Commands 203
7. Ad Replication 204
Ad Delivery and Replication 204
Server-Side Caching 205
What exactly is a tag? 206
8.Making a bigger impact with Display Campaigns 207

Chapter 3 - Section 3 - Rich Media 208
What is Rich Media? 209
Why use Rich Media? 209
Core Formats 211
Expandable ads 212
Figure 3.3.1a 212
Floating ads 213
Figure 3.3.1b 213
Video ads 214
Figure 3.3.1c 214
Pushdown and Sidekick 215
Figure 3.3.1d 215
Data Capture 216
Home Page Takeovers, Wallpaper and Page Skins 217

Figure 3.3.1e 217
Figure 3.3.1f 218
Polling 219
State Based ads 219
Local Connection ads or Sync Ads 220
Figure 3.3.1g 220
Print ads 221
Standardization and Publisher certification 222
From the top: Building Rich Media ads 223
1. Training the creative developer 223
2. Two schools of Rich Media specs: Publisher and Ad server 224
3. Parent and Child files - the polite load 226
4. Expansions and Floating -the declaration of state 228
5. Interaction tracking: Default and Custom Events 229
6. Interactions as additional clicks 230
7. Building Custom Rich Media ads 231
Qaing and Testing Rich Media ads 232
Trafficking Rich Media ads 233
Publisher Troubleshooting 234
Rich Media for Mobile & Rich Media for Instream 235

Chapter 3 - Section 4 - DCO 237
Standard XML Display 239
Figure 3.4.1 240
Figure 3.4.2 241
DCO Design Phase 242
Creating the dynamic template 242
Figure 3.4.3 243
Dynamic Creative Planning 243
Figure 3.4.4 244

DCO - Ad Creation Phase 245
Single Versioning or Manual Creation 246
Excel Mass Versioning 247
Mass Versioning from a Feed 247
Figure 3.4.5 249
Using a web service to restrict load on Advertiser hosted images 249
Figure 3.4.6 251
DCO - Assignment to Placement Phase 253
Placement vs. audiences 253
Figure 3.4.7 254
Targeting different Versions to different Audiences 255
Figure 3.4.8 256
Connecting DCO with Audience Buying 257
DCO - Optimization Phase 257
Auto-Optimize Vs. Manual Ad Controls 259
Figure 3.4.9 259
Auto-optimization as a "Black Box Technology" 259
Multi-variant testing 260

Chapter 3 - Section 5 - Instream Video 261
What is Instream? 262
Figure 3.5.1 262
Pre Roll Mid Roll Post Roll 263
Figure 3.5.2 264
Site Served Video 265
Figure 3.5.3 265
Where do Video Ad Networks add value? 266
VAST - the video tag standard 267
Companion Ads 267
Figure 3.5.4 268
Ad Skipping and TrueView 268
VPAID - Standardized Interactive Video 269

Figure 3.5.5 269
VPAID Authoring 270
Figure 3.5.6 271
DCO for Instream 271
Figure 3.5.7 272
Beyond VAST and VPAID 273
Publisher Certification 274
Figure 3.5.8 276
Trafficking Instream ads 277

Chapter 3 - Section 6 - Mobile 279
The Connectivity Consideration 280
Internet Enabled Device Divergence 281
Publisher Certification for Mobile 282
Figure 3.6.1 283
Challenges within the browser 283
Responsive Design 284
Whatever happened to Flash? 284
The fallback Mechanism 285
Trafficking ads for Mobile Browsers 286
QA for mobile browser-based ads 286
Figure 3.6.2 287
Challenges within the app 287
Figure 3.6.3 288
Trafficking ads for In-app ad campaigns 289
Where do Mobile Ad Networks add value? 289

Chapter 3 - Section 7 - Paid Search and Natural Search 291
SEM and PPC 293
Creating Campaigns and Keywords in the Engine 294
Figure 3.7 295
Figure 3.7.1 296

Connecting a Bid Management System to the Engine 297
Figure 3.7.2 298
Why track paid search through a third party ad server? 299
How to track Paid Search through a third party ad server? 299
Manual Search Tracking 299
Figure 3.7.3 300
Automated Search Tracking 300
Figure 3.7.4 301
Ad Server tracking integrated with Bid Management 303
Figure 3.7.5 304
Extending tracked Paid Search to the long tail 305
Figure 3.7.6 306
Extending tracked Paid Search to Mobile 306
Search and other Channels 307
What is SEO? 308
Figure 3.7.7 309
Referral URL to trigger a click 310
Figure 3.7.8 312
Ad Serving costs for Search 313

Chapter 3 - Section 8 - Affiliates and Affiliate Networks 315
Affiliate Sites and Networks 315
How the Affiliate Earns 316
Auditing the quality of the Long Tail 317
Figure 3.8.1 318
Tracking Affiliate Activity in third party ad serving 318
Figure 3.8.2 319
The absence of impression tracking 320
Affiliates that grow 320

Chapter 3 - Section 9 - Email, in-game and Social 323
Permission Marketing 324
Figure 3.9.1 325
Sending Emails on Mass 325
Email Marketing advertising metrics 326
Third party ad server tracking for Email 327
In-game and digital outdoor tracking 328
Social Media tracking 329
Figure 3.9.2 330
The versatility of tracking pixels 331

4 **Chapter 4 - Section 1 - Conversion Tags and Attribution** 334
What are conversions? 335
Figure 4.1 337
Where to place the conversion tag? 338
Figure 4.1.1 339
Where to have conversion tags (table) 340
CPA 341
Conversion tag types - Sales Vs. Counter Tags 342
Figure 4.1.2 343
Extended data/Extended Parameters 343
Figure 4.1.3 344
Figure 4.1.4 345
How do you test conversions? 347
A. Conversion Tag Functionality - Piggybacking 350
Container Tags 350
Certified Container Tag providers 351
Which conversion tag fires first? 352
Selective Pinging 352
Figure 4.1.5 353
Piggybacking with Extended data 353
Figure 4.1.6 354
B. Conversion Tag Functionality - Conversion Windows 355

Window Settings 355
Post-impression and Post-click windows 356
Long and Short Conversion Windows 358
Shadow Tagging 359
Reporting on Conversions 360
Repeat Conversion Windows 360
C. Conversion Tag Functionality - Attribution Modeling 360
Figure 4.1.7 361
Attribution Settings 362
Engagement Mapping 363
Figure 4.1.8 365
Custom Attribution Modeling 366
Challenges in implementing conversion tags 367
Tag Management Systems 368
Figure 4.1.9 369
Universal Tagging Security 369

Chapter 4 - Section 2 - Audiences and targeting 373
Audiences and Targeting at a media/inventory level 374
Audiences and Targeting at a creative level 375
Third party ad server targeting rules 375
Layered Targeting 377
Figure 4.2.1 378
Geo-targeting 379
Behavioral targeting 381
Figure 4.2.2 381
Contextual targeting (page-scraping) 381
Figure 4.2.3 383
Search keyword targeting 383
Targeting based on Verification Data 384
Additional Data targeting 385
A/B Testing 386

Figure 4.2.4 388
Tru lift 389
Figure 4.2.5 390
Audience Assignment 391
Channels which bend the rules 392

Section C - Ad Server Analytics 395

5 **Chapter 5 - Reporting Technology** 397
What is ad server reporting? 398
Aggregated Data 399
Figure 5.0.1 399
Figure 5.0.2 400
Reporting API 401
Third party Ad Serving vs. Site Side Analytics 401
Figure 5.0.3 402
When is the Data Processed? 403
Real time reporting 403
Processing Clock 404
Where is the data processed? 405
How is the data filtered? 405
Where is the data displayed? 406

Chapter 5 - Section 1 - Basic Aggregated Reporting 407
QA Reports 408
Global reporting 408
Figure 5.1.1 409
Delivery reports 410
Verification Reports 410
Performance Reports (Conversion and Attribution Reports) 410
Cost and ROI reports 411
404 error reporting 411

Expiring cost packages 412
Channel Specific Reports 412
The Standard Metrics 412
Impressions 412
A new impression? 413
Figure 5.1.2 414
Clicks and Click-through Rate (CTR) 414
Conversions & Conversion rate 415

Chapter 5 - Section 2 - Ad Verification & Viewability 419
Effective and Ineffective Impressions 420
Reactive, Preventative and Pre-Emptive Verification 420
Figure 5.2.1 422
Publisher and placement Quality indicators 423
Contextual Verification Quality Indicators 424
Content Classification 424
Brand Safety 426
Figure 5.2.2 426
Demographic or Audience Verification 427
The panel-based Audience Verification Method 427
Figure 5.2.3 428
The social profile-based Audience Verification Method 429
Figure 5.2.4 430
Discoverability Verification Quality Indicators 431
Geo Reporting 431
Device Verification 432
Viewable impressions 433
Figure 5.2.5 434
Double Implementations and Fraudulent Serving techniques 436
Figure 5.2.6 437

Page Level Verification Quality Indicators 438
Domain URL Reporting 438
Other uses for Page Level information 439
White Lists and Black Lists 439
Figure 5.2.7 440
See through, View through and Recordability Rates 441
Figure 5.2.8 442
Preventative - Ad blocking 443
Pre-emptive - Verification in Programmatic 443
The Verification landscape 444

Chapter 5 - Section 3 - Advanced Aggregated Reporting 447
Cross Channel analysis 448
Day Part Analysis 448
Cross Channel Reports 449
Figure 5.3.1 450
Advanced Conversion Analysis 451
Time to Convert 451
Figure 5.3.2 452
Engagement Mapping Factors report 453
Reach and Frequency Reporting 453
Unique Reach 453
Average Frequency 454
Figure 5.1.3 455
Effective Frequency 455
Figure 5.1.4 457
Reach and Frequency Report 457
Figure 5.1.5 458
Adjusted Unique 459
Figure 5.1.6 459
Advanced Site and Creative Analysis 460
Unique Site Overlap 460
Figure 5.1.7 460

Ad format reporting 461
Global Benchmarks 462
Advertiser Specific Benchmarks 463
Figure 5.1.8 464

Chapter 5 - Section 4 - Channel Specific Reporting 465
Display Channel Reporting 466
Rich Media Channel Reporting 467
DCO Channel Reporting 468
Multi variant testing 468
Dangers in Version Overwriting 469
Version reporting 469
Figure 5.4.1 470
Instream Channel Reporting and Metrics 471
Video Started and unique video starts 471
Ad duration 471
Increment ad play-through (25, 50, 75, 100%) 472
Average % video completion 472
VPAID custom interactions 472
Mobile Channel Reporting 472
Cookie Duplication 473
Cookie Absence 474
Device Finger Printing 475
Cookies in applications 475
Device and Carrier reporting 476
Paid Search Channel Reporting 476
Search Overview 477
Keyword Pathway Analysis 477
Natural Search Reporting 477
Double counting and Single Counting 478
Figure 5.4.2 480
Affiliate Channel Reporting 481

Chapter 5 - Section 5 - Cookie Level Data 483
Reporting at the Cookie Level 484
Figure 5.5.1 485
How is the data supplied? 486
Look-back Data Dump 486
Cookie Level Feed 486
Report delivery, size estimation and storage considerations 487
Metrics in Cookie Level Data 487
Tracking impulse conversions 488
Pathway to Conversion Reporting 488
Figure 5.5.2 489
Point Solutions for Advanced Cookie Level Attribution 490
Reporting on Conversion tag Actions 491
Extended data at the Cookie Level 492
Figure 5.5.3 492
Figure 5.5.4 494
Impression and Click based token data at the Cookie Level 494
What is missing from ad server reporting? 495

6 **Chapter 6 - Optimization** 498
Questioning reporting data 499
1. Where did the report come from? 500
2. What kind of report is it? 500
3. Do the numbers look right? Does the data add up? 500
4. What time was the report pulled, what time zone is the account in and are these the correct date ranges? 501
5. Are these the correct metrics, conversion tags and hierarchy entities? 502
6. Do the names of the entities in the report match those in the campaign management interface? 502

7. What is the methodology behind these metrics? 503
8. Is the shown granularity providing enough or too much information? 504
9. Is there enough data in the report or were the traffic volumes large enough to gauge decent statistical accuracy? 504
10. Is there a chance that the data in the report be adjusted again later on? 505
11. Are all the rows and columns displaying? 506
12. Was the campaign setup in way that might affect a governed outcome? 507
13. Did the ad server encounter any kind of counting fault? 507
What is Optimization? 508
Optimizing to Campaign Objectives 508
Optimizing towards Delivery Goals (ideal for Brand Advertisers) 509
Unique Reach 510
Target Audience 512
Verification (quality safety) goals and impact 53
Optimizations towards Performance Goals (ideal for Direct Response Advertisers) 514
Engagement and Interaction 515
Total Clicks 517
Total Conversions or E-Conversions 519
ROI (& extended data) goals 522
Optimization Sample 523
Recording changes in Optimization strategy 523

7 **Chapter 7 - Retargeting and Privacy** 525
Retargeting Tags 526

Figure 7.1 527
Cookie Pools 528
Retargeting Windows 529
Retargeting for Creative Optimization 529
Figure 7.2 530
Dynamic Retargeting 530
Figure 7.3 532
Carousel Ads 533
Figure 7.4 534
Privacy and Behavioral targeting 534
The Ad Choices Program 535
Figure 7.5 535
Do Not Track 536
Retargeting for Media Optimization 538

8 **Chapter 8 - Programmatic** 539
Advertising Exchanges 541
Supply Side Platforms (SSPs) for the Publishers 542
Demand Side Platforms (DSPs) for the Advertiser 543
Figure 8.1 543
Connecting the ad server to the DSP 544
Figure 8.2 544
Quality Inventory Sources 545
The Second Highest Bid Auction 545
Real Time Bidding (RTB) 546
Cost Reporting in the ad server 547
DSP Strategies 547
DSP Retargeting with DCO 548
Figure 8.3 549
DSP Keyword Targeting 550
Negative Retargeting 550
DSP delivery controls 550
Building Private Exchanges 551

Figure 8.4 551
Strategies based on Publisher Audience Descriptors 551
Cookie Level Data Providers 552
Data Management Platforms (DMP) 552
Figure 8.5 553
Page Level Data Providers 554
Channel-Specific Complications in RTB 554
DSP - ad server of the Future? 555

9 **Chapter 9 - The Big Picture** 557
Figure 9.1 - Display and Programmatic 558
Figure 9.2 - Layering in Paid Search 559
Figure 9.3 - Layering in Affiliates 559
Figure 9.4 - Layering in Users, Advertiser Properties, Ads and Trafficking 560
Figure 9.5 - Layering in Retargeting 561
Figure 9.6 - Third Party Ad Serving - the Diagram 562

Closing Chapter 563

Further Reading and Study 567

Acknowledgements 569

Glossary 571

PURPOSE

Why write a book about third party ad serving?

Working in online media or digital advertising, there are certain levels of technical expertise which exist across any connected organisations in the space. From a distance it would be easy to assume that everyone has the same level of training but when it comes to using the technical tools which enable to work that we do, the space is complex and the idea of learning everything is daunting. Some of us that work in Online believe we do not need to learn everything, that the task is too great and the necessity, trivial. The truth is that the knowledge in any online advertising organisation is unfairly balanced between those that take the time to learn as much as they can and those that rely on the learned. Learning the role of the Media Agency planner, The Advertiser marketer, The Creative Agency developer or the trafficker only lifts the lid on a corner of the complex technical workflow of which we are all mandatory parts.

I wrote a book about ad serving for one reason: there are not enough educational resources about Digital Marketing, Online Advertising and third party ad serving to satisfy any of its stakeholders.

When a topic this complex comes along and the work requires an academic background which does not exist, the number of experts in the field shrinks and the circulation of verifiable knowledge becomes thin. The knock on effect is a culture of assumptions, a culture of the fear to ask too many technical questions of the few that have the knowledge and an environment which is risky to conduct business within.

This book provides in academia, a record of the technology as it stands today. I have grown up with technology that is growing at the same pace, and I am sharing what I know. I have tried my very best to not make any assumptions about how this technology works but despite asking many sources, the sources are too few to be completely verifiable across the board and for this reason I apologise for any incorrect assumptions I may have made about how this technology works. My aim is to keep the information up to date by offering revised editions and maintaining an internet presence; I am open to technical questions you may have regarding this material.

I currently work for a well-known third party ad server but I have tried my best not to be bias towards its technology. Indeed there are some features mentioned in this book that the technology provider I work for does not offer today but that exist elsewhere in the industry.

I really hope you get enjoyment and value from this book and that it helps your personal and professional understanding of these technologies. I put a lot of effort into bringing the parts together in a digestible way (which given the matrix-nature of the subject matter, is not easy). I welcome with open arms any feedback you may have regarding the content and any opinions I may have expressed within (my contributing editors have helped me to keep these to a minimum).

Thanks for buying and happy reading!

Part A

Setting the Stage

Part A
Setting the Stage

- The flow of this book
- What is third party ad serving?
- Getting some perspective
    - Users
    - Publisher
    - Agencies
    - The Advertiser
- Definitions and Methodology

The flow of this book

The interconnected nature of Online Advertising technologies reminds me of neurons in the brain. Their connected nature could be drawn two dimensionally and the resulting image would be a weave of connections. But the truth is that almost all nodes of this image would communicate with each other in some way, spanning distances that would make a visual connection in three dimensional space, too painful on the eyes to observe and too detailed for the brain to absorb. I was asked by colleagues and my network contacts as I was planning this book as to how I would lay it out given this almost quantum nature of the technology landscape. I decided to focus on the campaign management workflow and build concepts up from there.

In this sense the book starts with some background for the reader (Part A), relating each of the concepts back to third party ad serving, such as the base technologies on a user's machine and in "the Cloud" (Chapter 1) and the industry players or stakeholders that assemble and pay each other to bring an online ad campaign to life (Chapter 2). In Part B we look at campaign management

broken down by the various channels available today (Chapter 3) before looking at the concept of conversion tracking (attributing a sale online as a result of an ad campaign). Part C then focusses on understanding the analytic and reporting capabilities of an ad server (Chapter 5) before looking at strategies to act on that data such as Optimization (Chapter 6) Retargeting (Chapter 7) and Programmatic buying (Chapter 8).

This book also contains a whole heap of words and phrases that have been more or less adopted across the industry. Due to the infancy of the digital advertising industry the meanings of some of these words are not wholly universal or understood therefore there is a convenient glossary located at the end of this book for your reference. Words which appear in quotation marks can be referenced in the glossary at the back.

There are also some technology and company names that will appear in *italics*. These names are used as examples of technologies that may not be available in every market around the world but are growing or well adopted technologies used by big Agencies and Advertisers in Europe and the US. To provide the most balanced understanding of the technology landscape I have tried to offer the names of as many technology providers as possible although a large number of technology companies such as *Google* now offer many of the discussed features as part of their core *DoubleClick* marketing suite. The names featured are suggested examples. It is recommended that if you choose to utilise these technologies for digital marketing activity that you test their products first, preferably against a control group where a technology is not in use and also against a competitor tool which claims to offer the same results. This on-going experimentation will ensure that you are using the best tools available in the market today.

As the flow of the various chapters suggests, the campaign management workflow is cyclical; the setup of all the ads can be

changed at any point during a campaign based on trends in the analytics data which is accessible just by logging into the ad serving interface. In this way the ad serving interface is made up of two core areas: campaign management and reporting.

Perhaps most importantly let's address the elephant in the room:

What is third party ad serving?
A third party ad server is a business to business machine or system that does a specific service. This machine lives in various locations around the world to complete the tasks of the service; it lives in "the Cloud". As with other services that live in the Cloud it is possible to access them only via a web browser. Third party ad serving core "customers" are Media Agencies and Advertisers. These customers pay to have access to be able to log into the third party ad serving interface. Once inside the interface, the tasks are broken down into two separate areas: campaign management and reporting.

Customers choose to use a third party ad server for these main reasons:

1. Accredited Independent third party
2. Consistent Counting Methodology
3. Reliability to serve and Server uptime
4. Format Support and Publisher certification
5. Single location Reporting
6. De-Duplicated Reporting
7. Campaign Specific Needs and Features

Customers of third party ad servers pay initially for the ad space and then pay a tiny fee on top for each advert to benefit from the above. The third party ad serving space, as a section of the advertising technology industry has matured, the third party ad servers available to potential customers today is rather small. In

fact there are only really two truly global third party ad servers left choose from: *DoubleClick* and *DG/Mediamind*. The others have seen diminishing growth or have never seen adoption beyond a handful of countries and a global ad server is important so that global Advertisers are seeing the above benefits in all markets. These more local third party ad servers include *Adform, Mediaplex, AdRiver, Atlas, Flashtalking* and *Weborama*. Thankfully there is uniformity among these technologies - they all work in basically the same way so all the concepts in this book are valid across the board although one or two features may not be available in some ad servers.

Getting Some Perspective

Now with the global third party ad server humming away in the cloud, its tasks permeate out across the web affecting various groups of people. Let's now look in detail at who these people are:

Users

Everyone is a user, we all "use" the internet. Anyone accessing the internet or browsing the web is considered a user. This concept seems simple but if you work in online advertising you are a user but you also fall into another of the categories that follow. When talking to other media professionals about technology it always helps to declare your position: "In this problem that I am seeing with this ad; I am taking on the role of the user". Users will interact with the third party ad server when loading pretty much every web page and the total number of users in the world is around a couple of billion.

Publishers

Own a website with ads on it? Then you are a Publisher (your website is your publication). If you own or represent lots of websites you are a Publisher network. Publishers tend to see the world from their own perspective; they do not look beyond their own networks and use every tactic possible to drive "media

spend" (advertising dollars) into their own pocket. Advertisers and Agencies have tough jobs to do to spread that spend around effectively and this is where ad serving comes into its own. Publishers interacting with the third party ad server by receiving code from it show the third party ads on its pages and will have to work with its reporting systems. How many Publishers are there in the world: upwards of hundreds of thousands of big sites but in actual fact hundreds of millions due to the number of web sites on the internet.

Agencies

Agencies fit into two buckets. First there is the Creative Agency; who undertake the work of creating and building the ad (and do not need to trouble themselves too much with the world of third party ad serving). And secondly there is the Media Agency - a middleman in the advertising campaign workflow, managing and optimizing to make sure advertising dollars are spent effectively. The Media Agency has a tough job swimming through a minefield of technologies to achieve this. Any single piece of kit that can track all the ad serving data, automate planning and buying, costing and reporting are the magic bullet for a successful agency. How do Agencies use third party ad servers? The Creative Agency builds ads for the ad server and uploads them onto the system. The Media Agency will do the campaign management and reporting.

The Advertiser

The Advertiser is any company with products or services to sell that has enough money for a marketing budget for advertising. References in this book to the Advertiser are references to the marketing team working for that Advertiser. To afford to utilise enterprise level technology like an ad server there have to be a fair few media dollars in the bank for the marketer to use, yet well over ten thousand Advertisers globally utilise these crucial technologies. Few Advertisers choose to stand alone without a Media Agency and given the complex nature of the digital

environment it is no surprise. How does the Advertiser use a third party ad server? For Campaign management and reporting.

Selecting the right ad server is a tricky business, there are so many features to consider and yet the differences between them are very small. This is all thanks to industry standardisation which has been crucial to allowing Agencies and Advertisers to switch from one Ad server to another. It all begins with definitions and methodology.

Definitions and Methodology

The *Internet Advertising Bureau* (*IAB*); arguably the most widely recognised regularity body globally for digital marketing activity, has been instrumental in the shaping of digital advertising. Standardisation of words and phrases is probably the most important and fundamental concept in digital advertising. And although the *IAB* has done a lot of work to reduce miscommunication, there are still occurrences of the same concepts being given different names. For example *DoubleClick* have a "Floodlight tag", Atlas have a "Universal Action Tag", Mediamind has a "Container tag" (all are the same thing). The methodology that underpins what is meant by a "click" or an "impression" still causes headaches industry wide despite clear guidelines being set down. In online tracking it always comes back to methodology.

In my first job in the industry my line manager sat me down to read through thoroughly the methodologies of the big research companies; *Neilsen* and *ComScore*. These respected platforms were revealed to me as building their datasets on a sample of people from a panel. This meant that the reported figures in these systems were not accurate when low numbers were listed as they represented only one or two people and were not an absolutely accurate representation when scaled up (despite the growth of these panels in recent years, there is still some truth to these statements). Collecting information based on a panel of users as

opposed to measuring the behaviour of all users helps to demonstrate the importance of methodology; two datasets could differ wildly if one comes from a sample and the other comes from a full audience measurement. The methodology provides an indispensable context.

The lesson is that the devil is in the detail and technology companies are not always wholly transparent in their methodologies so it is vital to ask questions about any data. Sometimes two systems measuring the same things at the same time will not report exactly the same figures. The *IAB* has done its best to close the gap (such as allowing for a maximum of a ten per cent discrepancy between technologies measuring the same thing) but at the core these technologies are designed differently and might use different methods to "ping" servers at slightly different times, resulting in mismatched figures. This universal acknowledgment (like the unstable world of measurement in quantum mechanics) does not mean that all systems are flawed, but that using a single system to track all advertising activity reduces confusion and discrepancy because all things are measured under the same methodology. This requirement led to the advent of the tools we call third party ad servers (one counting tool used to measure all activity). Without getting too complex, it can also be argued that using multiple tools to measure the same things and then taking an average is the most accurate way to capture data. In an ideal world many Agencies would choose to use multiple third party ad servers for a single campaign but the costs often do not justify the rewards gained from a more accurate measurement.

Now that we have set the stage and given a very high level view of who uses third party ad servers and why they are used, we will now move to look at the neighbouring technologies on the web page and on the user's machine. These have an impact on how ad serving works and what the limitations are in data collection.

Chapter 1

Local vs. Remote Technologies

Chapter 1 - Section 1
Browsers and Devices

- The role of the browser in Digital advertising
    - *Figure 1.1*
    - *Figure 1.2*
- Tracking across multiple browsers
- first party Vs. third Party Cookies
    - *Figure 1.3*
- Cookie Deletion
- Distinction between Devices and Computers
- Ad blocking Tools
- Why is there advertising on the Internet?

Third party ad servers do not work isolation. Many technologies work hand in hand to form websites and web channels that create a user's internet experience (such web channels include "Search Marketing" and "Rich Media"). The ways these channels operate and grow affect the usability of the internet such as through web page loading speed and the pace of innovations in web technology development. This section looks at the tools and technologies that we come into contact everyday as internet users, (whether we are aware of them or not) and looks and how they each relate to the world of Advertising and tracking. The bottom-up approach to considering the capabilities of underlying tracking technologies (learning how core code like HTML and JavaScript relate to the capabilities of the browser) will provide a solid framework of understanding which can be applied to complex problem solving in ad campaign setup and optimization (e.g. if you understand how JavaScript works you can understand how to use it to adapt ad technology to fit unique problems and situations which will arise every day). Perhaps more importantly, this introduction provides a stable grounding from which to understand more about the web channels themselves.

Despite being such a ubiquitous technology the concept of the "Internet" is commonly misunderstood as just meaning the World Wide Web (the bit of internet that loads into the browser window). The Internet is in fact a far bigger entity and "the internet" in the context of online advertising means; the visual web, its applications and access points. Where information can be shared across the world automatically, this is the internet we speak of. It is important to be clear on this so that all industry stakeholders have a common understanding. When work with new clients, Advertisers and Publishers, the syntax and definition of such principles is best bedded in early on bearing in mind that the media industry attracts folk from all walks of life.

So as not to get our heads too far into the cloud (pardon the pun) let's just think simply about the browser to start off with:

The role of the browser in Digital advertising

Despite its paramount importance to the world of Digital Advertising, the browser-owning technology companies almost never engage in an active dialogue with any of the other Digital Stakeholders (Publishers, Advertisers and Agencies).The work of updating a browser and releasing a new version (however wild the impact may be on the design of Websites and Ads) has virtually no consideration or input from the vast Advertising community that it supports. This statement is therefore true:

> Browsers were not built to support the creative range of designs available for online ads and websites. Truly creative websites and ads have been predominantly born out of the skill and vision of artists and creative designers, who take the foundation of a new browser release and learn how to build the best user experience on top of it. The browser release will clearly provide new avenues to take code and design but the browser does not arrive with a flurry of supported ad and website designs.

The key takeaway here is not to ever expect support for any campaign from the browser themselves no matter how large the campaign may be. In my experience finding a technical contact working browser-side (even if you work somewhere like *Microsoft* and you are absolutely desperate to see if your campaign can utilise some kind of exception) is a goose chase not worth taking up and it's safe to assume you'll never be able to point out the *Mozilla* team in a crowd and buy these guys a drink. The browser remains an unknowable force working by its own rules, which the whole industry has learned to adapt to over time.

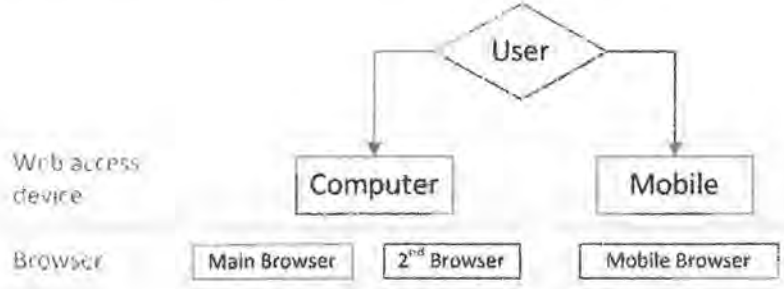


Figure 1.1: *Users access the web through an internet enabled device (Computer or Mobile) and then navigate web pages from within the browser.*

User opinion is divided on the best browsers to use to surf the web and the battle for the top browser is set to continue for many years to come. Firefox, Internet Explorer, Chrome and Safari dominate the browser space. But what relevance does the browser have to third party ad serving? (And does it make any difference which browser the user uses?).

The browser is very important to online Advertising because the user experience with the vast majority of the adverts on the internet is had inside the browser environment. Ideally the user

would be sitting at the same computer while using the web all of the time, as they surf the web and encounter relevant, innovative ads. In fact to store and track the most accurate data in the third party ad server, a single user would need to browse the web on a single device (using the same browser all the time). In this single browser the user would rarely delete their browser cookies and this browser would have acted as the single conduit for the full record of that user's fully web journey, engaging with ads and using the internet. The user's web activities are occurring in just one place (one browser and one device).

We are now going to start referencing "cookies". In simplified terms, a cookie can be thought of as a document. On this document is written a unique number which identifies the user called a cookieID. Now, if each user had only one cookieID and the cookie that it belonged to was never deleted, every user would have just one unique number. If you were to count all the total unique cookieIDs in existence, you would have an accurate reading of the total number of users on the Internet. As you may have gathered, things are not quite as simple as this.

First if all, cookies are not stored on a computer or device in a folder called "cookies". Each browser has its own "cookie-vault" which does not let other browsers in. So if one did want to compare and contrast users between browsers, the process required to do it is somewhat complex.

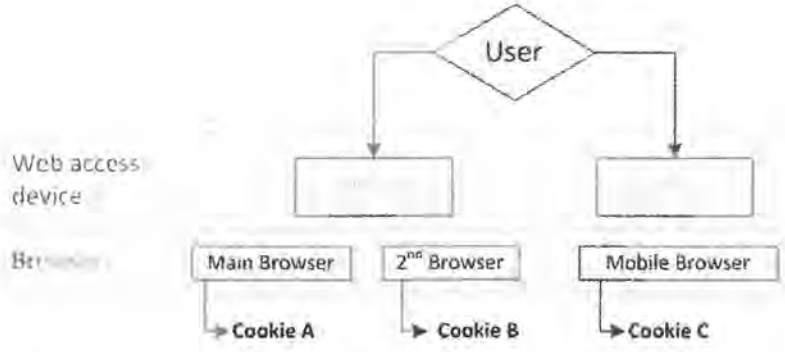


Figure 1.2: *Each browser has the facility to store cookies. Although a cookie is meant to represent a single User, the use of multiple browsers creates multiple cookies - a known industry problem.*

Tracking across multiple browsers

The second and perhaps more obvious problem is that users very rarely use a single access point or single browser to use the internet. Instead users come on from different devices and browsers, sometimes using more than one device at the same time. This is a pain for online advertising for the simple reason that multiple cookies look like multiple individuals in the data that comes out of the ad server after it has completed tracking and processing its gargantuan datasets.

Users who use multiple browsers frequently are harder to track as activity is recorded against two separate cookieID's which are never matched and technically speaking the ad server sees two separate individuals.

This is a danger for the marketer and the Agency because the true number of users seeing adverts is skewed. If there was just a single browser in existence then online ad measurement would be far more reliable. In the reporting chapter later on in this book we will see that ad server reporting still does not offer a solution to rooting out the uniqueness of all cookies when faced with the problem of users that use multiple browsers. Bearing this in mind

when conducting analysis of the ad serving data may help in providing a clearer picture of the total number of users interacting with ads in the campaign. Let's look at cookies in more detail.

First party vs. third Party Cookies

There are two main types of cookie. The "first party cookie" is placed on the user's computer or 'dropped' by the website the user is visiting. By dropping, we mean that a remote machine or server sends a tiny file to user's computer or device, whereby the file is delivered and remains until it is deleted or removed. This file is a cookie and usually contains the cookieID, the timestamp of when the cookie was dropped and the filename for the file will reference the domain it came from. If a website drops a cookie file onto your computer from its own domain, it is a first party cookie (The web address of the website is part of the websites "domain" so in http://autos.msn.com "msn.com" is the domain and "autos." is the subdomain). If a website drops a cookie on your machine from another domain (such as the domain belonging to a third party ad server) this is called a "third party cookie".

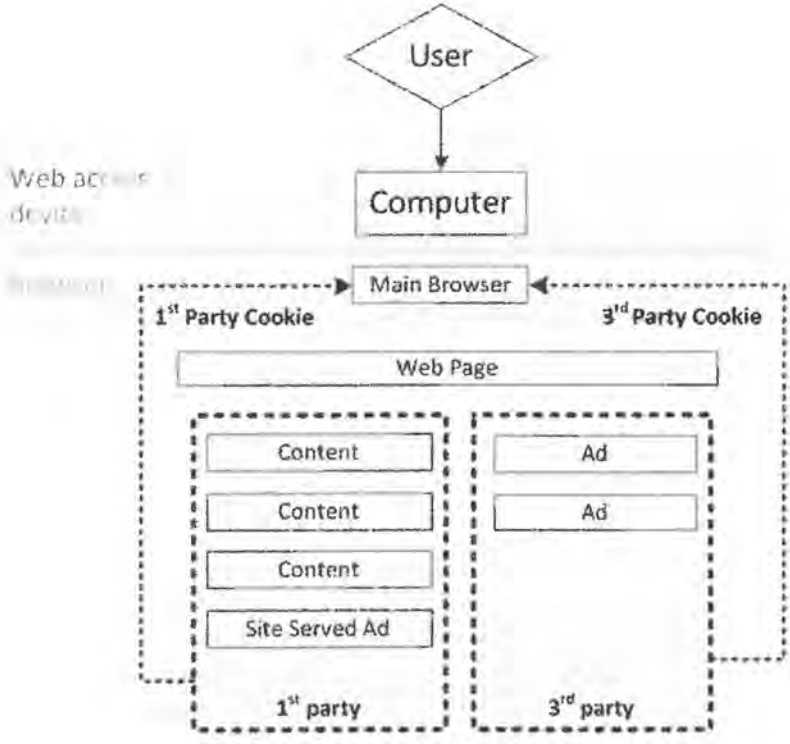


Figure 1.3: *Where a user uses a single browser, browsing to web pages triggers the downloading of content and ads. The native identifiers for the user belonging to the content and then ads are first and then third party cookies.*

Cookies are a website's way of remembering some information about you as a user because once they have been dropped they remain on the machine until they are removed. Since the cookie is residing on the user's machine it is said to 'live' on the "client side" and the cookie remains as a signpost or unique reference point, for a website to lookup that unique number against its own database (where there may be more information about each cookie stored on the databases our its own machines; this is called the "server side")

This lookup or matching process allows a website to 'remember' things like some settings or changes the user may have made to a website. Most commonly the first party cookie can be used to remember the information the user put into forms on a website, (so that when they visit that site again, they save time by not having to re-input this information). This is very useful when you, as a user, have to login to a website each time and it remembers your username for instance.

Third party cookies such as those belonging to the third party ad server have a record of the user's activity matched alongside engagement with an ad campaign. The ad server can report on whether the user saw an ad, clicked on an ad, or went on to buy a product after seeing or clicking an ad online. To be clear, it is not the actual user's name that the server is remembering, but the unique reference number stored in that user's cookie. The server stores no personally identifiable information and so never knows that you are you; it just knows that your cookie allows the server to remember behavioural information about that cookieID. Those readers who live and breathe *Excel* and databases can begin to picture a master table sitting in the deep backend of the ad server, where every ad engagement can be tied back to a cookieID (see Chapter 5 - Cookie Level Data)

Cookie Deletion

A new consideration if you were just beginning to get comfortable with cookies and cookieID's:

1. Cookies do not live forever.
2. The more time passes the higher the likelihood of a user deleting a cookie.

Cookie Deletion affects third party ad serving because each new cookie is treated as a new user. Users that delete their cookies frequently skew the data available in the ad server since there is

no way to tie a deleted cookie from a user machine to the new cookie dropped by the server to replace it. As we will see later this affects a lot of different reports including Unique Reach and Frequency and Performance reports (see Chapter 5).

There is lots of research posted around the internet regarding cookie deletion. Some technologies claim to be able to forecast and predict cookie deletion in different counties. Some ad servers offer alternative reporting methods that can account for cookie deletion (See Chapter 5 - Reporting - Adjusted Unique reporting).

It is in the best interest of Advertisers, Agencies and Publishers to prevent cookies from being deleted, the benefits of their use are hotly debated in the forum of rights protection organisations in the field of personal privacy and it is here that the Digital Advertising stakeholders defend the use of cookies where information that is held about users does not lead to a direct identification of that user. In this legal minefield the debate continues globally as to the acceptance and regulation of cookie-use, where the user is given the option to opt-in or out of the automatic action of cookie dropping (see Chapter 7 - Retargeting).

Distinction between Devices and Computers

The web browser on a user's phone or tablet offers a different experience than the one on a desktop computer. Examples include the screen size and mouse-free navigation. Considerations should be made for these differences when planning an ad campaign, from the building of the ads to the design of the website that the user should end up on when clicking on the ad. In some cases, Advertisers will pay a Creative Agency to produce different types of ads and design different websites for users using different browsers.

This variety increases the cost of running an ad campaign, and in some instances is used to preserve the same user experience across devices and in others offers a creative diversity. Ad serving

technology allows for browser-specific targeting (see Chapter 2 - Publishers - Publisher Ad Servers) to enable different ads to be displayed on different devices.

Technically speaking "mobile" browsers are just different versions of the same pc and mac based browsers that users the world over are now used to using. As with desktop device browsers, the differences between the versions can affect the way that ads and pages appear or are rendered. This presents a real problem for consistent ad delivery, especially for more complex formats like Rich Media (for definition see Chapter 3 - Rich Media). The Advertiser or Agency may plan with the intent of offering the same user experience to all users across all devices but to do so means testing ads on all such devices before a campaign begins. The same is true in the world of Publishing where all web pages must be checked in all new devices and browsers too. Some Agencies have opted to fill vast rooms with all of the available consumer internet-enabled devices for the purposes of such testing. Others will try to emulate the multitude of available browsers using software found online.

HTML5 may finally be a light at the end of a long awaited tunnel, as it is a leap forward in standardisation for both page and ad rendering. HTML5 has been devised to try and eradicate the need for device by device testing, which has extended the testing time required before a campaign launches, creating industry-wide inefficiency.

HTML5 is one of the most important things to happen to advertising in recent memory since the timelines of an ad campaign will affect all parties; and the improvement in the flexibility in the coding of an ad can mean the difference between a frustrating internet experience and a fully-fledged engagement between user and ad. HTML5 is not really a component of the browser itself but a mark-up language, something we will look at in more detail in the next section. Beyond HTML support,

browsers bring plug-in technologies to consideration and these are controlled by the user. One plug-in type in particular has historically caused a stir in the birth and growth of Online Advertising and these are Ad blocking plugins.

Ad blocking tools in the browser

Some users install software or browser extensions (also called "plugins") to "block" ads from appearing on their machine. These tools disrupt the delivery of the image of the ad or in fact any information from the Publisher-side ad server (See Chapter 2 - Publishers) to the user's browser. Ad blocking was more functional in the days of dial-up when bandwidth was precious and loading an ad meant an extra ten minutes of page load time. Today bandwidth is so quick that ad blocking would serve very little purpose in speeding up page load times (ads usually make up a tiny percentage of the "weight" of the whole page as the files are downloaded from the server). Instead these ad blocking tools are used to provide an ad-free internet experience. All stakeholders operating in Online Advertising need to be able to respond to questions about Ad Blocking but since the mechanism is so simple it is the wider impact of these tools that requires focus as an Ad-Free internet experience is a threat to the whole medium of the web. This in turn concludes our summary of the relevant features of the browser in Online Advertising.

Why is there advertising on the Internet?

With such far reaching potential, the web comes with its own politics and the politics frame the development and demands for further supporting technologies. The web is predominantly used, by its users to source and digest written and (more recently) video content. Content takes time to produce. I once set about building a content heavy website: it took hours, days and weeks of time to compile. Apart from learning obtaining new skills and having the warming experience of contributing to a giant database of useful information about the world, I was only compelled to continue adding such content by receiving money in return.

The two most compelling options for a website owner to produce quality content are to charge a subscription fee for the service (plaguing the user with a cost out of their own wallet) or to seem to offer the content for free to the user, at the cost to the Advertiser (who are really just paying for the attention span of a "targeted" user).

Ad blocking is not an undetectable weapon on consumer combat; the failed attempt by the server to deliver ads is detected each time it occurs. As this failure rate grows, there is an inverse decline in the size of the audience for the website/Publisher to offer to their Advertiser base. A decline in the audience size means a decline in revenue, which ultimately leads to a decline in both the amount (and more importantly) the quality of the content because less money is available to help produce it.

Content is king online. The best content draws the most attention and although the growth of user generated content (UGC) has been rapid and far reaching; the quality user "traffic" remains with big Publishers where it always has done (and more niche "long tail" audiences hop to more obscure websites with lower levers of traffic). Thankfully technologies like Advertising Exchanges are making some ground in balancing the books, but this is a topic for Chapter 8.

Let's move on from the browser to the content of the web; the websites and the ads themselves and the base technologies on which they are constructed.

Chapter 1 - Section 2
The Anatomy of Sites and Ads

- What the user sees
  Figure 1.4
- HTML
  - *Figure 1.5*
- JavaScript
- Images
- Text Links
  - *Figure 1.6*
- Flash
- The advent of HTML5

What the user sees

Take a magnifying glass and peer into the source code of any web page (or just use the browser buttons to load up the page source) and inside you will find a rather difficult to read set of instructions and code. The educated eye can pick out commands in the code that help identify how the web page has been constructed. This includes information about where the ads come from as well. Both HTML and JavaScript are the most common languages to find inside this source code. They load the content in the browser window, rendering as a website constructed by a set of common elements or entities. Text, Images and Flash make up the vast majority of the web today (soon to be overtaken by an explosion in video). Without these components the web could not support advertising as it exists today. Understanding more about them personifies their importance, highlighting technical encounters on all sides of the Advertising ecosystem.

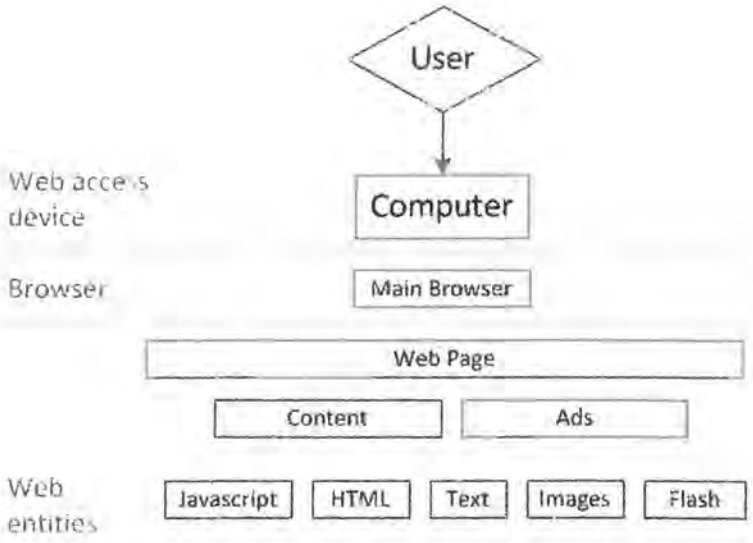


Figure 1.4: *Navigating to the Web Page via the main browser, the user encounters Content and Ads. These are made up of what can be classed as entities: technologies and code that form the very foundations of the familiar web.*

HTML

"HTML" (Hyper-Text Mark-up Language) is the core 'code' of the web based browsing experience. It literally is the code that describes to the browser what the content of each web page should be and the browser renders this code to give the effect of a page of information. HTML is the language of "tags" and when we refer to "tags" in third party ad serving, we are talking about a few lines of HTML code. This code can be embedded directly onto a web page by a web development team working on the website and when a page is loaded up in a browser by a user the code 'executes', and an ad or more content is displayed. (To see an HTML ad tag, skip to Chapter 3 - Display Ad serving Operations - Ad Tags).

HTML is how almost all of the ads served today are requested from the ad server. If the environment in which an ad needs to be served does not support HTML, then other code can be used to make that request (such as in a computer gaming environment). In fact in most instances, the HTML is used to house a more complex request which occurs in JavaScript. JavaScript does the fetching, carrying and calculating of the content. HTML is simply a way of carrying the required ad code that the Publisher can quite easily slip onto their websites.

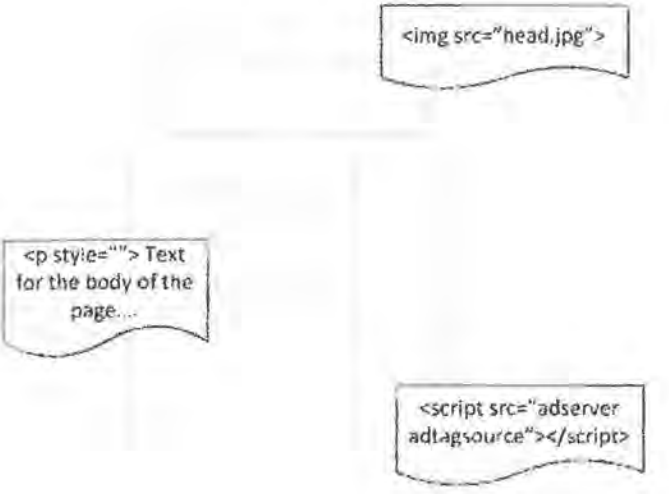


Figure 1.5: *Behind the scenes a standard web page is made up of a series of HTML tags. An Ad tag is no different, which can sit on any web page to display an ad to the user. In this instance (bottom right) the HTML tag contains a reference to JavaScript.*

JavaScript

JavaScript is called a client-side language. It runs or "executes" in the user's browser and allows for a whole host of important web page tricks such as popup windows and browser manipulation (like a command to print something on the page). JavaScript can

gather and send a user's information and to a remote server so that the server is capable of responding by sending out a new piece of information (like making a previously hidden part of the website, suddenly appear).

JavaScript is the most common method of ad delivery that third party ad servers use. The flexibility of the JavaScript code allows the ad to be pulled directly into a Publisher's web page from the third party ad server and make decisions about what ad to display in real time. JavaScript can do interesting things like work out what time it is, or find out what browser the user is using: information like this can be passed back into the ad server when the ad loads allowing for some impressive customisation inside the ad itself. It is not common for an Advertiser to want to get this technical when running a campaign but every now and then JavaScript does come to the rescue when an ad campaign just wants to be a little cleverer or more innovative.

To understand how JavaScript is essential in the technical process of ad request and delivery we will look in more detail at the relationship between Ad Serving and JavaScript in Chapter 1 - Basics on Server hardware. For the time being it is sufficient to say that JavaScript returns one of four remote entities from the ad server and pulls them into the web page as the page loads. These four entities are Image files, Flash files, text links and video files. Video files are a little more complex to explain so this is explained after providing a more thorough technical introduction - details for video can be found in Chapter 3 - Instream Video).

Images

Images transformed the early web from a completely text-based environment to something that could be digested as an entertainment form by the masses and advertising was one of the first places to utilise the ability of early browsers to load images. The GIF type of image was more favoured by the advertising space because GIFs can contain more than one image at a time

and so a very rudimentary form of frame by frame animation was possible. The trouble with GIFs is that they are very restricting in terms of creativity. They are not interactive and images look poor next to Flash. Today, GIFs are still used in online advertising as a "fall-back option" in case the user's browser cannot support Flash (after all, a grainy image of a company logo is better than completely wasted and white ad space). The traditional *.gif file extension is also no longer mandatory and better image compression methods such as in the *.png format or *.jpeg formats are commonly utilised. The legacy of the ad as an image remains with ad serving technology today where all standard display ads are comprised of an image portion of the HTML code as well as a click or text link portion. The two sections which make up this HTML ad tag allow all display ads to report back to the ad server; counting the loading of the image as an "impression" metric and the user action of the "click" as a click metric (more detail on this in Chapter 5 - Reporting).

Text-Links

Since HTML code allows any piece of text, object or image to be clickable by the user via the use of the "Anchor tag", a simple text link can be issued by the Ad Server as the most basic type of Ad. These text-link ads are extremely common all over the web and are mostly utilised by ads that exist in Search engines. By issuing an image and text link together, the image is clickable but the two parts of the code can be separated out so that the image can simply be programmed to appear without the ability to click, or the text link can be implemented without any image present.

1. <!-- Begin Ad Tag code -->
2.
1. Starts here
4. <img src="">
5.
6. <!--End Ad tag code -->

Figure 1.6: *A sample of a very basic non-JavaScript ad tag. The numbered rows are for reference purposes: Row 2 describes the opening of the anchor tag: the URL described is the clickthrough URL. The code in rows 3 and 4 is surrounded by the anchor tag which acts as the click redirect portion of the ad tag. Row 4 can contain a 1x1 otherwise known as a tracking pixel.*

The text in **figure 1.6** (row 3) is clickable inside the web page that the user can see. This is the "text-link". Text-links are the most basic building block of an online ad; their use and behaviour are described in a series of widely recognised phrases. As many of these sound similar it is important to clarify this with participating parties when setting up and running an ad campaign.

- The action of a user clicking and being directed to an Advertiser site is called a "**click through**".
- The process of a user completing the clickthrough successfully is called a "**click-to-land**".
- The destination address is called a "**click through URL**"
- The portion of the ad tag which declares the click through URL as part of any ad or text link is called a "**click redirect**".
- A click redirect may also be referred to as a "**click tracker**"
- The equivalent for measuring the loading of the ad (even as a text link) is called a "**tracking pixel**". (Covered off in more detail in Chapter 3 - Display Ad serving).

So far we have described the two elements of an ad tag as being the image portion and the click portion. However it is not common for a "static image" (i.e. a still picture) to be used as the visual ad; rather the image displays as a Flash animation.

Flash

Having a familiarity with Flash is essential to working in the Display and Rich Media channels in Digital Advertising. "Flash" are the animated and interactive panels dotted around the

internet that can house video, get bigger, or show multiple images in an impactful delivery.

Flash is nascent to the web browsing experience and Flash files or ads themselves are created in *Adobe's* Flash authoring environment (more on this in Chapter 2 - Creative Agencies). Flash is so widely used and adopted across the web because of its ease to create animations and its flexibility to be creative with the code and the development environment. In Chapter 2 we will look in much more granular detail to understand more about Flash and its capabilities. Flash is so popular that it seems that it will be in use across the internet forever but new technologies are emerging that will replace it. HTML5 is the favourite.

The Advent of HTML5

HTML5 is created to amalgamate all the capabilities of both Flash and HTML in an experience which provides a consistent user experience on the web. Flash was popular for so long because of a similar consistent usability but the flaw in Flash is the need for users to have installed the Flash player plugin. With more devices and browsers available to users, supporting the plugin has become difficult for *Adobe* and since using the Flash authoring tool comes at a cost to the creative developer, but the development of HTML5 is free, HTML wins as the standard for the future. With HTML5 there is no need for Flash, as video and animation become part of HTML. The vector animations and functions that Flash built fame for facilitating are now supported in HTML5 and the community of supporters for this new standard are stronger where the updates and maintenance of the authoring environment of HTML5 will become increasingly faster and stronger than Flash.

So far we covered off internet technologies to give context to online advertising. The next section moves focus from the local user computer to the servers.

Chapter 1 - Section 3
Ad Serving

- What does the ad server do?
    - *Figure 1.7*
    - Table of Definitions
    - *Figure 1.8*
- The role of the CDN
- Waiting around for Ad Delivery - the user experience
    - Pre-ad display
    - During ad display
    - On ad close

This section will touch lightly on the relationship between the user's computer and the technologies sending packets of information to that user's machine. This is useful to explain what is really going on beyond the familiar glare of your own PC and browser. The words and phrases contained in this section are commonly thrown around in technical conversations where an Advertising technology is involved.

What does the ad server do?
Many other available texts describe how user's machines and servers exchange information to load the various pages of the web. The focus for this section is to explain the exchanges that occur between machines to distribute, load and count ads right across the web. As a reminder the Third Party Ad Server has two core functions:

1. The function to serve or deliver the ads to the users machine
2. To count and track the delivery and performance of these ads

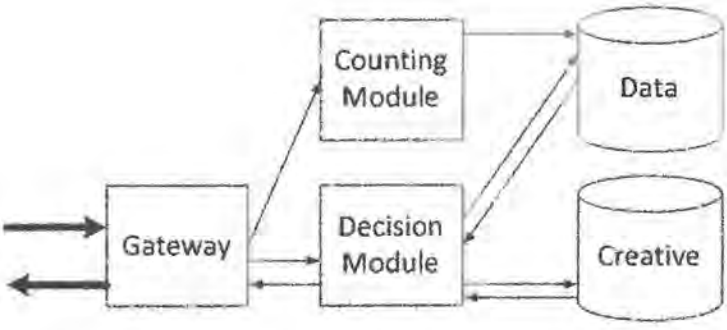


Figure 1.7: *Data in and out on-site at the ad server. This highly generalised diagram shows how the user request into the ad server gateway gets split for the two functions. First the ad server counts using a counting module and in parallel makes a decision to serve the correct asset back to the user.*

Figure 1.7 shows the fast function of the ad server to count and deliver an ad but the bold arrows on the left only show data coming into and going out of the ad server. The internet infrastructure that will carry this data to this point undergoes a longer journey. Before we undertake this journey we need to make sure we are using the right language to describe the various ports of call and processes that the data undergoes.

	Definition
Local Machine	The user's computer or device to access the web
Remote Machine	The server that the user's computer is exchanging data with
Host Server	Another name for remote machine, the server is "hosting" the users request for content or ads
Packet	Data is transmitted between machines in

	packets. Utilizing the best server route between the local and remote machine. Sometimes packets get lost and this can have a minor impact on the counting of ads served.
Router	To connect to the ISP the user's machine must utilize a local switch or router. Again, all data comes through here.
ISP	Internet Service Provider; the user's access point to the web, all data packets come and go through the ISP
CDN	Content Delivery Network; a large network of geographically located content servers allowing users to download content from a regional source. Ad servers themselves do not distribute replicas of the files all over the world. They work with CDNs to achieve this.
Data Center	Data Centers are the physical buildings that house ad servers and servers of the CDN. As the user's browser makes a call directly to the ad server to request an ad, the ad server must deliver the information back and count the request as soon as possible. Having strategically located buildings in global hubs and highly populated cities ensures a faster delivery.
Server Farm	A non-technical term to describe a data center or group of data centers
Ad Request	The initial loading of the ad tag on a Publisher page sends an HTTP request for to display the correct ad to the ad server. This

	is called an Ad Request.
JavaScript Response	The ad server responds to the ad request with a response; this is response predominantly occurs in JavaScript.
Impression Tracker	All ads contain Impression Tracker code which sends data as an "ad request" to the counting module or "ad counter" at the ad server
Impression	When an ad counter receives and responds to an HTTP request for a tracking asset from a client (*IAB.net* Guidelines)
Click Tracker	When the user initiates a click, this code routes data to the "ad counter" at the ad server to count a "click"
Ad Call	Another name for Ad Request
Ad Play	An industry wide disputed meaning, not backed by the IAB.
Ad Close	When a JavaScript served ad is closed or the browser leaves the page, the Ad can make a final request to the server

Before we look at the bigger picture it is important to note that the technical delivery of the ad is not the same as an industry agreed count of the impression metric (which can be found here http://www.iab.net/media/file/Global_meas_guidelines.pdf). However as this definition has certified industry adoption, it is taken as to the true measure of delivered ads across all ad counting systems.

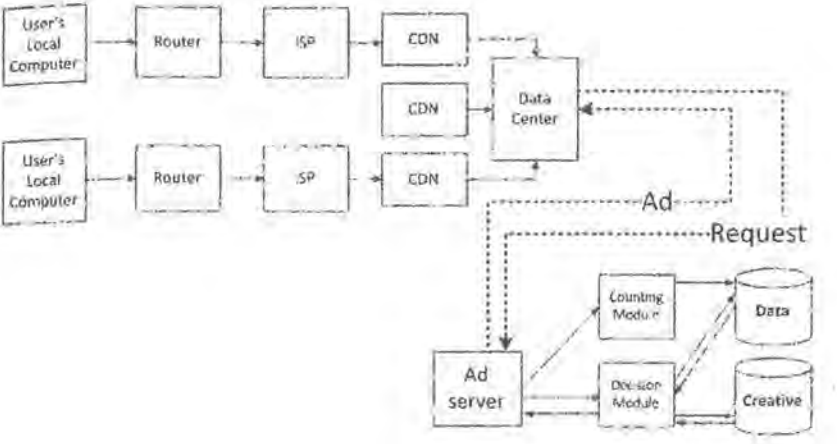


Figure 1.8: *User requests channel up to the ad server's data center where counting and ad serving decisions are made. As multiple requests are being made by multiple users for the same assets, the ad server decision module does not need to distribute the same ad every time but rather lets the CDN distribute the ad.*

The role of the CDN

The Content Delivery Network traces all the flight paths in the world, joining up the dots between the cities and the ISPs (Internet Service Providers) around the world and routing replicated information onto users' devices where-ever they may be. Ads stored in the ad server need to be replicated in the ad server data centres, prepared for the most strategic delivery of the assets (a delay called "ad replication") and a process that occurs when the campaign is being "trafficked" (more on trafficking in Chapter 3). When an ad call is received the ad server transmits the ads to the CDN's. *Akamai*, described by some as the internet's backbone, is one of the largest CDN's and ad servers utilise its services consistently. Larger Publishers like *MSN*, rely on *Microsoft's* own technology rather than outsource to systems like *Akamai*. The *MSN* CDN has the capacity to allow other internet

technologies requiring CDN distribution to utilise it for such needs.

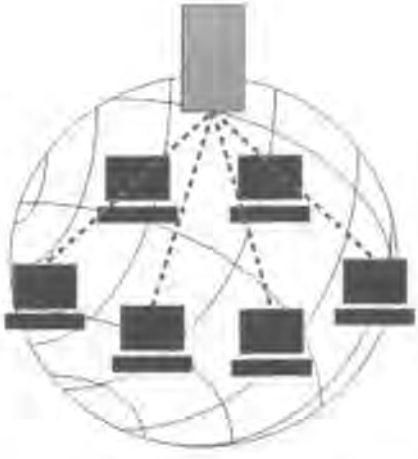


***Figure 1.9**: CDN content delivery is most frequently represented in common conversation by reference to "The cloud". By offering servers in local proximity to user devices internet access speeds are greatly increased and files can easily be replicated across the world, this includes both site content and ads.*

Waiting around for ad Delivery - the user experience

So what is the effect on the user? Is it glaringly obvious that the most miniscule of delays actually has a shed-load of processes behind it? For the most part the answer is no, but with so many processes in place, once in a blue moon, the user will be stuck on a page where an ad fails to load. The purpose of the ad servers reporting capability is to quantify the investment clients make in online advertising, including demonstrating consistent ad delivery and performance which the Agency or Advertiser will reference frequently to ensure their messaging is reaching its audience, if not only to check that there is nothing clogging up the pipes (more details in Chapter 5 - Reporting). Providing there are no problems the optimum user experience will look like this:

Pre-ad display

1. User arrives at Publisher page
2. Ad request is sent to Publisher ad server, while the request for the content goes to the Publisher CMS
3. The Publisher content loads
4. The Publisher ad server calls the third party ad server
5. The ad request arrives at the local data center depending on the users location
6. The counting module counts an impression
7. The decision module selects the most suitable ad
8. The ad server sends a JavaScript to initiate the loading of the ad
9. The JavaScript loads on the Publisher page and loads in makes calls to all the necessary assets
10. The assets are pulled in from the CDN which in turn pulls in data from the ad server creative database
11. An Image or Flash execution appears in the user's browser window

During ad display

12. The ad server counts an "Ad Play" - which is an Instream Video metric (See Chapter 3 - Instream) or delivered impression metric
13. Piggybacked third party impression trackers load (See Chapter 3 - Standard Display)
14. Viewability begins counting (See Chapter 5 - Ad Verification)
15. Rich Media constantly reports "interactions" to the counting module (See Chapter 3 - Rich Media).

On ad close

16. Viewability ends counting and sends data to the ad server counting module (See Chapter 5 - Ad Verification)
17. Rich Media interactions that utilise a timer are reported to the ad server counting module (See Chapter 3 - Rich Media)

This exposes the mini cosmos of work that the ad server is doing. It is clear why the work is all done by machines in faraway places. The next section turns attention back to the online advertising industry to look at who the players in this space, and how many people and hands it takes to launch an online advertising campaign and make the most of ad serving technology.

Chapter 2

The Industry Players

Chapter 2
The Industry Players

- The agency holding groups
- Who are the big Agencies?
- Running a Media Agency
- Creative Agencies
    - *Figure 2.1*
- Publishers

To frame the context of advertising technology the books include information about the various industry players. Given the nature of the advertising industry to shift clients and accounts, this section may well fall out of date quite quickly; however this is an attempt to represent it at a high level.

The agency holding groups

The largest and most powerful advertising Agencies sit under holding groups which themselves are home to a multiverse of Creative Agencies, PR Agencies and Media Agencies. With so many Agencies under a single umbrella, careful precautions are taken to ensure that Agencies can take on the work and business of many competing Advertisers without allowing sensitive information to spread to those competitors or separated divisions. Since a holding group may house ten or more Media Agencies which themselves may house ten or more Advertisers and somewhere down the line two Advertisers are bound to fall into the same vertical. In some instances Media Agencies and holding groups have been known to build out whole new Agencie just to ensure that such Advertiser data is kept very separate. The power of the holding group tends to fall with the price it can get for publication ad space, the economies of scale of such enterprises drives down the price of ad space, with the promise c business on an enormous scale.

Who are the big Agencies?
There are plenty of Up-to-date sources online that can give a more recent and thorough coverage of the breakdown of the Agency groups. As a very high level overview:

- **Publicis Omnicom group** houses the media Agencies formerly of the Omnicom group: OMD, OMG, PHD and Manning Gollieb and the media Agencies formerly of the Publicis group such as Starcom Mediavest and Zenith Optimedia.

- **Dentsu/Aegis** brings to the table Agencies such as Iprospect, Carat and Vizeum.

- **WPP** offers the industry GroupM and its heavyweight's Mindshare, Mediacom, Maxus and MEC to name but a few.

- **Interpublic group** (commonly shortened to IPG) includes Universal McCann and Initiative

- **Havas group** houses MPG/MediaContacts and Havas

Running a Media Agency
Media Agencies are typically made up of account teams (an account being an Advertiser). Each team will be made up of planners, buyers and account managers. Media Agencies have supplier contracts with technology providers including third party technology providers. These agreements vary considerably and can differ between region and country. Media Agencies also have contractual agreements with Publishers that outline payment terms. Such terms are usually confirmed in writing before each advertising campaign via the Insertion Order (IO) a document which also clarifies the amount of media space purchased as well as conditions of the campaign. Larger Media Agencies are occasionally affiliated with creative Agencies or more often have

their own in-house creative developers and designers. Much mo common is a contractual relationship between the Advertiser an the Creative Agency directly. Media Agencies usually work with Advertisers in a long term capacity by pitching for their business by responding to and winning the chance to respond to a brief issued by the Advertiser. From an online perspective Media Agencies take a big burden off of the shoulders of the Advertiser. The work required to plan and buy media to reach the Advertiser's stated campaign objectives requires many hands an larger advertisers will rarely fill out their marketing teams with the headcount to undertake these tasks. It is for this reason that third party ad serving technology is most utilised by Media Agencies, who report their successes to the hiring Advertiser.

Creative agencies

There are an enormous number of creative agencies and the landscape is very competitive. A creative agency will usually produce a mock-up of how a campaign could look with lots of creative ideas, when pitching to an Advertiser to build a relationship with them and attain their business. Advertisers usually have preferred Creative Agencies that they will utilise for building creative across various media channels including TV and Print. The majority of Creative Agencies are small companies employing just a few designers and developers but some have reached enterprise sizes and fall under the same holding groups as the Media Agencies.



Figure 2.1: *A single Agency holding group may have clients operating in the same vertical but the teams are kept separate, often in completely different buildings. The Advertiser themselves*

has the freedom to select a Creative Agency which might work in a completely different holding group or operate independently.

Publishers

Publishers are the content providers. They are web sites that attract a large number of users. To give up your day job and just live off of the advertising revenue from a website your site would need to attract in excess of one hundred thousand "unique users" per month (depending on the quality of these users: also called "the quality of the traffic"). This means that if the users are coming to the Publisher website for a purpose, and the ads shown are relevant and well-targeted to those users, then the quality of those users is more valuable. This quality traffic raises the price of the ad space on the Publisher website and would mean the Publisher can make a tidy profit from the money that Advertisers would be willing to spend to have ads on a Publisher website. Big Publishers (such as *Facebook, Google, MSN* and *Yahoo!*) attract a vast number of users and in turn are almost always included on Advertiser "Media Plans". Smaller Publishers, although not necessarily attracting big crowds have the added benefits of either a more niche audience or cheaper ad space. Publishers are the home of ads; they control the greatest cost to any ad campaign (the cost of the ad space) and at times view all technology in terms of their own properties.

The remainder of this chapter will be spent focussing each of these stakeholder groups. Each group spends more time using a different piece of ad technology therefore third party ad serving is more important to some and less to others.

Chapter 2 - Section 1
Advertisers and Marketers

- Who is the Advertiser?
    - The Marketer
    - Agency pitch
    - Relationship with the Media Agency
    - Relationship with the Creative Agency
    - *Figure 2.2.1*
- Controlling the page content
    - Controlling the pages of the Advertiser Site - the CMS
    - Landing Page
    - Basket Page
    - Confirmation Page
    - *Figure 2.2.2*
    - CRM Platform
    - Onsite Behavioural data
- Tracking the users on site
    - Site Side Analytics
    - Click to Land
    - Passing the tokens in the click-through URL
    - Referral

Who is the Advertiser?

Advertisers can be loosely divided into two buckets in the online world: "Brand" and "Direct response". Brand Advertisers tend to utilise the web to grow awareness and generate interest and engagement but rarely have something tangible to "sell" online. Direct Response (DR) Advertisers are usually much more complex in their advertising needs. They will have an e-commerce site to sell goods or services online and often need to produce ads and campaigns with objectives to target users at each stage in the purchase process. There are of course Advertisers that cross into both definitions. These are usually very large DR Advertisers who need a little of everything. In this section we will begin with an

overview of Advertisers, followed by their relationship with Media Agencies and then Creative Agencies. This is followed by an exploration around the predominant technologies an Advertiser will be working closely with; including the makeup of their e-commerce site and how to integrate such technologies effectively with third party ad serving.

The Marketer

Advertiser marketing team brings together the overall marketing strategy marrying up the available digital advertising budget with the right resource to get the best return.

For Brand advertisers, marketers are concerned initially with ad delivery and verification. If the ads have been served they should have been delivered at an optimal "reach and frequency" with the ultimate aim of hitting the targeting audience with a memorable creative message. In Chapter 5 we will look at how the ad server reporting can deliver this information.

For Direct response advertisers, marketers are concerned with ROI and want to demonstrate like for like or incremental sales from a set marketing budget. With the right combination of delivery techniques, reporting and optimization; digital is a very fertile ground for DR marketers with an ever increasing budget, to invest in.

For the marketer to ensure the advertising budget is best spent, the campaign goals are clearly laid out before a campaign begins.

Agency Pitch

To choose an Agency the Advertiser releases an RFP (Request for Proposal) to a shortlist of Media Agencies. This RFP details the advertiser's marketing objectives over the next few years and asks the Agencies to come back with the best strategic pitch for the work. The process can take quite a long time as the Advertiser is usually looking to sign a contract with an Agency over a term of a

few years; departing from such contracts can be very costly to the Advertiser, so the right decision is imperative. Once an Agency is secured, the relationship can begin to grow.

Relationship with the Media Agency

Once an Agency has won over the Advertiser the Agency account team begins the process of transitioning the "campaign workflow" process (the operational practices an Advertiser is familiar using) to align with the practices of the Agency or the needs of the Advertiser, where the Agency sees fit. This stage can involve moving the Advertiser from one set of technologies to another depending on the conditions of the contract with the Agency. This process can also take several weeks if not months but, done correctly, sets the stage for a blooming relationship in future campaigns. In the next section we will look in closer detail at campaign management, a process which is Agency-led.

As campaigns are processed the Advertiser and Media Agency learn which processes and practices work best for them and find ways of sharing and using the third party technology between them to suit everyone's needs. For instance some marketers prefer a more hands-on approach to the reporting data which the Media Agency is able to grant as long as expectations are clear about how decisions around campaign optimization are made.

Relationship with the Creative Agency

Unlike in the world of TV advertising, the Creative Agency takes a back seat to the role of the Media Agency in campaign setup (and usually the Advertiser will have a long term existing partnership with either a preferred Creative Agency or a myriad of them). The marketer will usually take a keen interest in the development of any creative to maintain a consistent brand image (the word "creative" is used to describe the state of a digital ad before it becomes an actual ad that can be used by the ad serving technology). This means there will be a back and forth of communication between the Creative Agency and the Advertiser

right up until the ad is fully developed, and it is mutually agreed that its visual elements will not change before the ad is seen by a public audience. Some big Brand Advertisers are more detail oriented than others when it comes to adhering to brand style guidelines and the sign-off process can get extremely lengthy.

For this reason, ad serving technology is designed to allow the Creative Agency to share a preview of the ads with the Advertiser directly and receive their approval before the ad is processed in the ad server for trafficking. The ad preview environment is particularly useful for Rich Media and DCO (See Chapter 3).

With its partners on board (Media Agency and Creative Agency) the stage is almost set to begin running an ad campaign. The final piece of the puzzle from the Advertiser's perspective is to ensure that users arrive at a website that extends their interest towards an eventual purchase.

Technical staff working on the Advertiser marketing team tend to place all their focus and attention on the Advertiser's site which can be pretty simple (for big brand advertisers that do not offer the chance to purchase online - perhaps having less than ten pages on their site) or vastly complex (such as in the case of a DR Advertiser like a retail catalogue website - where the number of pages to manage may be in the millions).

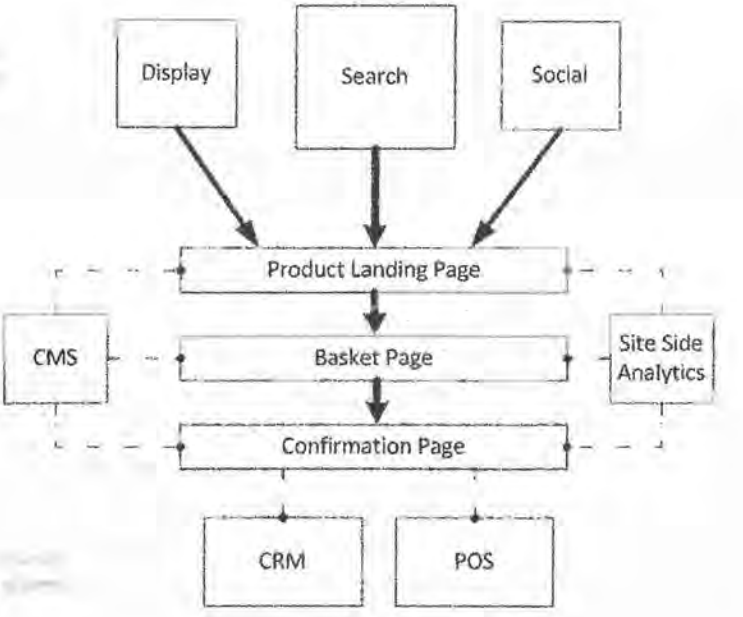


Figure 2.2.1: *In Advertiser-world, the Media Agency ensures that users get to the Advertiser site through various paid-for channels (tracked across the board only by the third party ad server). Once the user arrives on the Advertiser site there are several tools doing the measuring.*

Controlling the Pages of the Advertiser Site - (the CMS)

There are lots of types of pages on the Advertiser website but only a few of them are useful to the marketer. Pages which help demonstrate that a marketing objective has been fulfilled need to be carefully controlled, watched and measured. A CMS tool fulfils this first function.

An Advertiser website that has something to sell online is invariably going to contain a lot of individual web pages and a lot of information, to better inform the consumer on their journey to purchase. Web pages which are built in basic HTML code alone

are severely limiting. Thousands of live pages require a solution that will allow the web developer updating the Advertiser website to make changes in seconds. A Content Management System (CMS) allows for the creation and update of thousands of pages in seconds by providing a backend user interface for the web developer to change the site. The user of the website does not see the presence of a Content Management System but instead sees an illusion of "static" webpages. One of the crucial functions of the CMS is that it allows the marketing team to append the live pages with additional pieces of third party tracking technology to monitor page-level "conversion activities", these are called "Conversion Tags" and Chapter 4 is dedicated to explaining their purpose and uses in detail. Some marketers and Agencies need frequent access to the page code to change these tags but not all Advertiser sites (despite the presence of a CMS) allow for frequent changes. This can and often does directly impact the successful and timely completion of an online ad campaign.

In **Figure 2.2.1** the three core page types which best reflect the fulfilment of marketing objectives, are stated. Each of these page types is significant to the marketing effort and can be fully recognised as campaign yardsticks or "conversion events" in third party ad server reporting:

- The Landing Page (or Product Landing Page)
- The Basket Page (or pre-confirmation page)
- The Confirmation Page (or Thank-you page)

Landing Page

The Landing page is the first page that the user arrives at in the Advertiser's online domain. Being the first page, it needs to be designed to the user expectations. Thanks to the connected nature of today's search engines, a user can type in a keyword into the search bar, see a link to the Advertiser's site and click through to the landing page (more information about the Search Channel in Chapter 3). The landing page should therefore reflect

the very search term that the user inputted into the search engine and the Advertiser should customise the look and feel of the page to reflect assumptions about the user's reason for arriving on that page.

Basket Page

The user journey on an Advertiser's ecommerce website leads down, what marketers refer to as a "funnel", a lot of users reach a landing page, but many drop out en-route to the final purchase page. This analogy is often used in online marketing to describe the channel strategy for an online campaign. The consumer is attracted into the top of the "purchase funnel" with impactful brand ads that capture user interest, this is followed by capturing user intent and finally the user is drawn through the Advertiser site right up to final purchase at the end of the funnel.

The penultimate step is to load up the "shopping basket" with the intended purchase. Although the "shopping basket" element of the website itself may have many pages, the collective reference for one or more of these pages is a "basket page". Since a user can undertake various behaviours in the basket page, such as dropping out of the process, it is important for the Advertiser to monitor these pages with ad server tracking technology and make sure that the users are not lost forever (see Chapter 7).

Confirmation Page

The actual sale takes place somewhere between the basket page and the thank you page. The last key Advertiser page type is therefore the "thank you" page (which is also is referred to as the confirmation page); where the Advertiser confirms to the user that the product has been purchased.

Tracking activity on the confirmation page is crucial for DR Advertisers; a consumer only becomes a consumer once they arrive here, so the very arrival onto this page (which is not accessible by any other means other than by buying something) is

logged in the CRM and POS (as shown in **Figure 2.2.1**). These systems log a unique identification number related to the transaction called an OrderID. It is this same number which ties together the CRM, POS and ad serving data sets as they can all be setup to track the same value when an order is completed. The a serving data exposing the journey the user went on to arrive on the landing page and complete the sale - allowing the marketer t see exactly which customers respond to which ads and how much they spent.

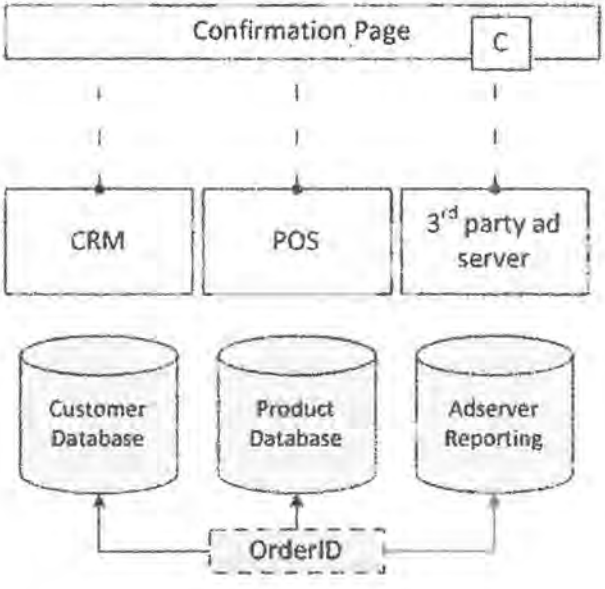


Figure 2.2.2: *The Advertiser can have the web development team working on the Advertiser site setup the Confirmation Page to pass the value of the OrderID field (declared on the page by JavaScript code) to the ad serving conversion tag. As the OrderID common to multiple systems it can be used to connect a user's behaviour together both on and off the site to determine how the decision to purchase is made.*

CRM Platform

In addition to product data, collecting customer data is vital to the success of the business. Customer data comes into the Custom Relationship Management tool from customer surveys and data collection points to the POS (Point of Sale) system. The data reveals seasonal buying habits, trends, price driven spending and it builds profiles of customers so that existing customers with certain buying habits can be modelled against new customers to predict buying behaviours.

Onsite Behavioural data

To enhance CRM databases, and add extra value to their businesses some Advertisers present opt-in tick boxes in their sign up processes which allow the data they collect, about those users on the site, to be used again for other purposes. Sometimes it is sold to third parties data providers to aggregate with other data sources before being sold on to other Advertisers. Marketers, who fail to do more than collect user details off the back of a campaign, could fund future campaigns by selling on the data. Of course a well-nourished CRM will pay dividends later on when answering crucial marketing questions.

From controlling the point of entry and the pathway through the site to the point of sale, marketers obtain data about the "order" itself but less so about the consumer behaviour in the confines of the site. The second platform that the marketer will spend lots of time in on the Advertiser side are the Advertiser or Site-Side analytics system, something that the marketer will have almost exclusive access to use.

Tracking the users on the Site - Site Side Analytics

Some technologies get more press coverage than others and "Site Side Analytics" tools are often confused with third party ad servers because they are one of the few technologies that have

their own entire shelves in large book stores. Established Site-Side Analytics providers include *IBM's Core Metrics*, *Adobe's Analytics Suite* (formerly called *Omniture*) and the popular, free alternative; *Google Analytics*. Site Side analytics technologies work by the Advertiser placing a snippet of code on every page of the Advertiser website and the tool measures the incoming and outgoing site traffic (mapping the journey users take through the site).

Click to Land

Since the first time the Site-Side analytics tool sees a user is when they first arrive at the landing page, it is common place for a marketer to try to match up the users clicking through on ads (as reported by the third party ad server) to those users landing on the landing page (as reported by the site side analytics tool). Unless the two systems are integrated this is a flawed exercise and will raise many discrepancies because the counting methodology between the two systems is inherently different.

Passing through Tokens in the click-through URL

Integration is possible where the click-through URL of the ad contains an additional variable that the site side analytics tool can recognise when the user arrives to the page. These variables are referred to as "tokens" and when setting up the campaign, the trafficker will append the click-through URL to either carry through a fixed or a dynamic token. To keep matters simple, let's focus on the fixed token scenario:

A click through URL might look like this:
http://www.advertiser.com/productpage1

A click-through URL with an appended token might look like this:

http://www.advertiser.com/productpage1?campaignname=campaign1

In this example the actual click-through URL was changed by the trafficker in the ad server interface (to see how this is done head to Chapter 3 - Ad Serving Operations). The site-side analytics tool could be setup to recognise the variable called "campaignname" and use the information passed through attached to this variable to get more information about the incoming user (such as the campaign name for the ad they just arrived to the page on). In this instance *campaign1* is a fixed value; it doesn't change unless the trafficker changes it. In the scenario of providing a dynamic token the ad server might be clever enough to have some backend "macros" or shortcut pieces of code that will dynamically add a value into a field.

In this instance the trafficker might insert:

http://www.advertiser.com/productpage1?campaignname=%campaignname%

The macro is shown with a percentage symbol either side, and when the user clicks on the ad, the clickthroughURL is formulated in real-time so that the macro-portion pulls in a master-value from the ad server backend called "campaignname". This feature is a huge time saver for the trafficker.

The trafficker can add:
?campaignname=%campaignname%
to all click-through URLs in multiple ad campaigns for the same Advertiser and the campaign name will dynamically change for the site side analytics tool to read the information and correctly attribute the user and their arrival to a campaign name which matches the master list of campaign names in the ad server.

Referral
Either way, the site-side analytics data makes for compelling analysis. Essentially these tools measure user journeys through the site from entry to purchase. The data is then used by the web

development to structure the website and by the marketing tea who use the data to enrich the CRM database. Site side analytics reveals important information about how users arrive at the website. This is known as "referral" information and can disseminate between traffic arriving from paid search vs. natura search. Where these tools fall short are in identifying user behaviour before the user arrived at the site which is where site side analytics reporting marries up nicely with third party ad serving data.

It is the role of the marketing team to match data from:

- The CRM
- The POS
- The third party ad server
- The Site-Side analytics tool

...to answer some or all of the following questions:

- Which ads are driving the highest return? And what is the optimum customer journey?
- What affect do campaign optimizations have on custome behaviour?
- What are the behaviours of the most valued customers and how can this data be used to create a bigger and mo loyal customer-base?

We have looked in detail about the world from the perspective o the Advertiser and Marketer. It is now time to turn our attention to the role played by the Media Agency.

Chapter 2 - Section 2
Agency Planning and Buying

- The role of the Planner
    - How to plan?
    - *Figure 2.3*
    - Research Tools
    - *Figure 2.4*
    - Demographic Digital Media Consumption by Publisher Website
    - Unique Audience - Site Overlap
    - Research through Benchmarking
    - In-Agency planning solutions
    - *Figure 2.5*
- Technology for Buying
    - Agency Trading Desks
    - *Figure 2.6*
    - The Agency-side Media Plan
- Selecting an Ad Server
- What technology-based Services do Media Agencies offer?
- Media Agency data analysts

Media Agencies bring a wealth of expertise to the table but not all Advertisers employ them to manage and run digital ad campaigns. Those Advertisers that feel they have the talent and buying power, utilise technologies and outsource some functions (such as trafficking) to achieve the same goals as hiring a Media Agency. The risk inherent in this approach is that the full scale of the work undertaken by the Media agency is cloaked by design as a positive customer experience so that Advertisers can spend media budget without getting their hands dirty with the day to day operations. The following section looks at the predominant roles and functions undertaken by the Media Agency.

The role of the media agency can be divided into four key areas, which we will look at in more detail:

- Planning Campaigns
- Buying Media space
- Managing Campaigns (locally, regionally and globally)
- Reporting and Optimizing to Advertiser KPIs

The role of the Planner

Deciding where to spend marketing dollars to meet Advertiser campaign goals, is a tough chore and takes considerable experience to get it right. Agencies employ "planners" for this very role. The process starts by having the planners dissect the campaign brief received from the Advertiser. They then decide which sites and channels will yield the best results.

How to plan?

The channel recipe for campaign effectiveness is not an exact science owing to the buying behaviours of the consumer. Strong cross channel planning strategies are balanced by the marketing principles of "the purchase funnel", suggesting that Display advertising is effective at the top of the funnel for consumer awareness and research, while the Search channel is more effective lower down the funnel where there is an intent to purchase or click. In Chapter 3 we look in depth at each channel its merits and how to manage campaigns via the ad server effectively.

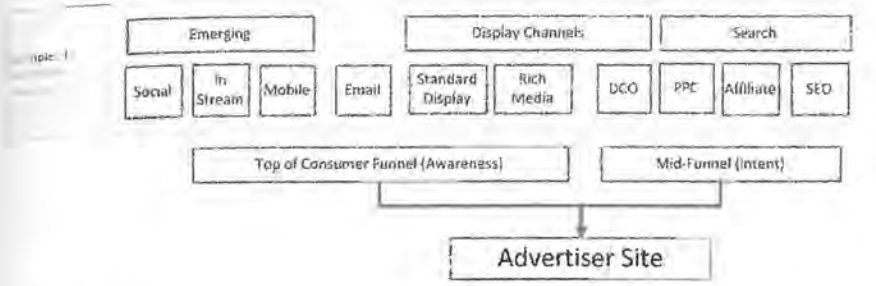


Figure 2.3: *A generalisation of the paid digital marketing channels available for the Media Agency to plan and buy ad space across. Broadly speaking the right mix can be formulated through careful research and succinct planning, although new channels and Publishers make the optimum mix a challenge to get right. More on Campaign Setup by Channel in Chapter 3.*

Research Tools

Planners and buyers use research tools to establish which sites in a preferred channel mix will meet the demographic targeting and audience reach criteria to satisfy the Advertiser objectives for the upcoming campaign. *Neilsen* and *ComScore* are the two dominant research technologies in digital advertising. These two companies closely track a panel of individual web users across the globe with their consent (identified as true individuals rather than anonymous cookies), and use this panel as a sample of the overall internet population. The sample is then scaled-up to the predicted total audience size of the internet (called the "internet universe") and the demographic details of the unique visitors present on each large Publisher web site.

The output could be compared to a league table of websites where comparisons can be made across sites with the greatest reach (called a "unique audience"), the count of users and the "page views" for each site (the number of times all users are

loading individual pages). Given that we are talking about samp sizes the numbers are not precise reflections of all users on the sites, but they do give a clear enough indication about Publishe website traffic for a planner to make some decisions. The plann will log into the research tool before the start of a campaign to decide that the audience they are aiming to target with a particular campaign brief; uses a certain Publisher website. The next step is to pass such information along to the buying team the media agency who instigates the purchase of available advertising space on that site.

Autos		A	B	C	D	E	F
		Nielsen			**ComScore**		
		18-25 m	25-40 m	60+ m	18-25 m	25-40 m	60+ m
1	Autos.microsoft.com	80%	70%	25%	90%	65%	30%
2	Autotrader.com	75%	75%	20%	80%	80%	20%
3	Webuyanycar.com	70%	40%	10%	75%	40%	15%
4	usedautos.net	50%	20%	5%	40%	20%	7%

Figure 2.4: *An agency that uses both the Neilsen and ComScore research tools can export reports from those platforms and compare the data for the two systems.*

Demographic Digital Media Consumption by Publisher Websi
In the example in **Figure 2.4** the top sites for a given time perio for websites classified in an automotive vertical show that for t first site on the list (microsoft.com) that 80/90% of the total audience of 18-25 year old males on either of the measured audience panels for the two technologies visited this site. If it i the intent of the Advertiser (that the media planner on the

Agency account team is servicing) to try to reach as many 18-25 year old males as possible, it would be wise for the planner to inform the media buyer to contact microsoft.com to try to secure ad space for an upcoming campaign. If the planner was trying to reach one hundred per cent of the 18-25 male audience, both of the research tools agree that it is not possible to do this with a single site.

Unique Audience - Site Overlap

To reach one hundred per cent of the right audience, still using **Figure 2.4** as an example, the Agency would need to consult a different report exported from these systems called a Unique Site Overlap report. Such a report would show the uniqueness of the audience on each site in the right demographic bucket. At this early planning stage the purchase of the media space that results from making decisions by looking at this data is not a one hundred per cent accurate reflection of what will actually occur when the campaign goes live. Only once the campaign is live and all users are tracked (as opposed to using a panel) will it be clear where there is overlap in unique users. The ad server can report on this directly (see Chapter 5) with its own version of the Unique Site Overlap report, but the actual delivery of the ads to the demographics stated in this table needs to be verified. *Neilsen* and *ComScore* offer additional tools based off of their panel data to verify the delivery of ads to the correct demographic group (see Chapter 5 - Verification)

Research through Benchmarking

Data collected in the third party ad server from previous campaigns can be extremely useful in determining the best Publisher sites to choose for planning as the historical data can be viewed in a way that allows for Publishers to be ranked by performance across all of the metrics available in ad server reporting. Big Agencies may amalgamate all the data for all the Advertisers they service to give a good indication of publisher performance by Advertiser vertical. Agencies that have run

enough campaigns historically can find out if a combination like the ad dimensions will give better performance on one site over another. Such tools might be developed in house by Agencies but some independent technology providers have alternatives. *Mediamind's Smart Planning* tool is an option for Agency planners looking to improve on a pre-campaign planning strategy. This ad-server owned research tool offers benchmarking across Publisher websites for a large pool of Advertisers, thereby giving a more accurate picture of where certain ads might perform effectively, long before the media space is bought and paid for.

In-Agency planning solutions

Informed decisions about the campaign plan, lead to conversations between the Agency buying team and the Publishers. The Agency buying process has evolved from its origins by traditional methods of buying media is still used alongside the new technology-supported buying methods. In traditional buying, the Agency buyer makes a call to a contact they have acquired at the Publisher site where they are interested in buying media space. On both sides of the phone, the Publisher and Agency agree on a particular "buying-rate" or "cost model" for the Advertiser's campaign. The Agency buyer may negotiate a position so that the Publisher is responsible for performance of the campaign and in doing so the final price of the ad space may be tied to this figure (more on actual cost models and costs in Chapter 2 - Publishers). The relationship between the media buyer and the Publisher seller in most instances will have been formed over many months in which time the Publisher would meet with the buyer to provide updates about the Publisher's media space and technology integrations. In instances where Publisher performance may not have reflected well in research tools like *Neilsen*, such relationships win the Publisher repeat business.

The request for ad space or media "inventory" from a Media Agency to a Publisher is recorded in an "insertion order" or IO

(essentially an invoice). The signing of the IO secures the deal and the campaign can begin.

Technologies have surfaced, created by either the Media Agency themselves, or by third parties to automate the setup of a campaign following the confirmation of a deal and keep a record of the negotiation proceedings. These are called "planning tools". The *Donovan Data System (DDS)* product called *iDesk*, is a good example of one such technology. *iDesk* allows the Agency to quickly input the details of the purchased ad space, connect this to their billing system and automatically "setup a campaign" inside the third party ad server interface. *iDesk* is currently setup to automate this process with all the major ad servers (see Chapter 3 - Display Ad serving for more information about how this works). In addition to this Agency buyers can come back to the records in this system in the next campaign or buying cycle and make sure that throughout negotiations that the Publisher does not attempt to charge a higher price than previously.

There aren't a huge range of options for Media Agencies that are looking to upgrade, replace or install a billing system that communicates well with planning tools and outputs data that allows an agency account team to keep an eye on the pacing for marketing budget spend across the campaign. The future promises to bring these kinds of budget management capabilities directly into the ad server (partly managed through the ability to input cost data in the ad server interface in the trafficking process as explained in Chapter 3) but for the time being the most widely adopted system for literally issuing payments to Publishers and keeping track of all costs: is *MediaBank* (also part of the *DDS* family under the brand name *Mediaocean*). This is crucial because the marketing budget must not just cover off the cost of the advertising space, it must also cover the cost of using technology like 3d party ad serving and additional Agency fees.



Figure 2.5: *the bread and butter of Media Agency world consists of Research, Planning, Buying and Trafficking in a straight-forward campaign development workflow. As the plans for the campaign begin to take hold and IO contracts are made with the Publisher team, the Media plan develops as a living document, referenced by multiple parties over the campaign's lifetime.*

This buying process is key to Media Agency involvement in digital advertising. The size of larger Media Agencies mean that on paper a large number of IO's will be double-signed with the largest Publishers. Buying more ad space affects the price as Publishers attract Media Agency dollars through economies of scale. At a certain volume commitment, prices are lower for media Agencies than they would be for Advertisers because Advertisers will almost always be buying less ad space.

Technology for Buying

Buying from Publishers and Publisher networks directly is now referred to as a "traditional media buy" because as we allured to before: technology has changed the buying process. Getting on the phone to secure ad space (inventory) from a Publisher leaves too much of a dependency on human contact, relationships and it takes too much time. Automatic buying technologies now allow for the purchase of mass reach inventory known as a Publisher

"remnant inventory" (literally the ad space which is unsold). The processing and placing of orders for inventory is now done through "ad exchanges" which are trading marketplaces for media inventory (not wholly dissimilar from the stock market) which can also occur in real time with the correct enabling technologies in place (RTB). The traditional media buying conversations still occur for the "biggest deals" such as large Publisher homepage takeovers; the purchase here is for ad spots with a higher value or larger audience known as "premium inventory".

Agency Trading Desks

There are a limited number of known and popular Ad exchanges where the inventory is grouped and classified by the Publisher by the audience expected to be exposed to that ad space. At the most basic level it can be said that the exchange advertises the sale of an audience. Accessing the Exchanges is hard work and the technology that allows the purchase of these ad spaces and audiences are DSPs (covered in Chapter 8). The "Agency Trading Desks" (ATD's) are specialist buying teams that are part of the Agency that utilise the DSPs to access the Exchanges. They have close ties to the Agency buying team and are utilised to purchase audiences on the Exchanges, a process undertaken when the traditional buying method is not used. This is a fairly rich topic so I am going to save the details for Chapter 8. The basic premise is that Agency buying teams don't just have to pick up the phone to Publishers anymore; they can buy the ad space through in-house technology teams which utilise the DSPs now. Big business agency trading desks include *WPP GroupM's Xaxis* and the *Publicis Audience on Demand* (*AOD*) platform.

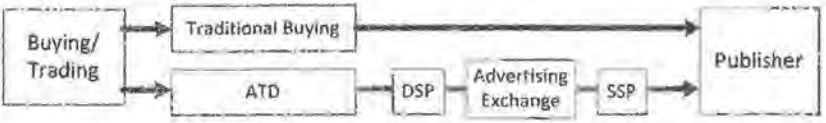


Figure 2.6: *The process of Media Agency buying and securing media inventory is possible down two pathways. The traditional method is a direct relationship-based selling opportunity. The ne methodology; predominantly used to trade remnant inventory utilises Advertising exchanges. Expertise in automated trading is managed by the Agency trading desk team. More on this model i Chapter 8.*

The Agency-side Media Plan

Once the buying process is complete the account team will have in their hands a copy of the "media plan". The media plan is a single document which clearly shows which ads are due to run and when. As the document is only completed when all parties are satisfied with the inventory purchase, it contains all the relevant information required to begin the construction of the ac themselves. This is due in part to the supply of the Publisher specifications (details of which can be found in Chapter 2 - Publishers) and will include information like ad dimensions and when the campaign is due to go live. Once complete, the media plan is supplied to the "trafficker" on the Media Agency side as well as the Creative Agency. The process followed by the Creative Agency is detailed in the next section. As for the trafficking process, the complexity of the proceeding processes is dictated b the make-up of the media plan. For instance the media channels that the campaign is planned and booked to run on will determin on what processes and considerations need to take place in the trafficking process (such as for trafficking ads due to run in *Googl* search versus those due to appear on *Facebook* or the *MSN* homepage. The nuisances of this channel by channel trafficking have been fully documented in the Chapter 3 – Ad serving Operations.

Once trafficking is complete, a campaign is active and the Media Agency account team will share the responsibility to monitor the delivery of the ads and report on campaign performance using the reporting functionality in the third party ad server. The specific reports that Media Agencies can see are demonstrated and explained in Chapter 5 - Reporting. Finally the Media Agency will complete the process of "campaign management" by either optimizing the ads in the existing campaign (See Chapter 6 - Optimization) or will utilise the information collected to run a second campaign or extend the existing campaign (See Chapter 7 - Retargeting).

The campaign setup process, from research right through to trafficking should be a carefully considered operation, but time constraints around the availability of inventory and the timelines of the Advertiser's own marketing plans can mean a rushed process. From speaking with Creative Agency, Publishers, Media Agencies and Advertisers involved in daily campaign management over the years I have managed to determine that any digital ad campaign needs at least two weeks to get the planning and setup just right. Rarely will a campaign get such a generous window and parties are often rushed on delivery. Every experienced trafficker can tell stories about 'Friday night' trafficking, a well-known burden of the trafficking role in which media plans for the weekend are received at the last working hour of the week, but a rite of passage for working in digital advertising.

During the stages of campaign setup and management (research, planning, buying, trafficking, reporting and optimization) the Agency account team will continually check in with Advertiser, bringing them up to speed on how close to those marketing KPIs, the campaign is. On some occasions the planners and buyers will manage to secure more marketing budget to deliver more ads on the media plan to reach a goal. The optimization process is also strongly connected to buying and Agency buyers may have to revise their IO's with Publishers if the ads are not performing.

This concludes the common roles of the Media Agency account team. As we have seen Media Agency's sometimes have trading desk teams as well. Quite common are also Trafficking and Technology teams. Such a team would be made up of Agency employees that are technical and skilled working within the environment of the third party ad server. Where these teams exist, the in-Agency trafficking process is slightly different and this team might take on the full process of trafficking a media plan instead for all the other account teams and their campaigns. Such teams also handle the troubleshooting of problematic campaigns, ads and provide advice around solving technical queries. It is also common for this team to conduct the selection process for choosing new technologies to use including ad servers.

Selecting an Ad Server

As this book will reveal, choosing an ad server is not an easy task. The selection process should be careful considered based on both the needs of the Agency and the needs of the Advertiser. As third party ad serving technology is rapidly evolving, traffic and technology teams on the Agency side will typically invite third party ad serving companies to pitch for new business (which tends to be more for on-boarding new Advertisers or migrating whole Advertisers from a different ad server, rather than going through the consideration process for each new campaign). After a period of research the traffic and technology team will issue an RFP which contains a list of technical questions about capabilities and higher level business questions about costs and plans for the ad server in question. If the RFP and the pitch that follows meet the Advertiser needs the Agency will formalise the relationship and complete an ad serving contract with the selected technology provider.

What technology-based Services do Media Agencies offer?
Media Agencies sometimes own technology solutions and sell these on to their Advertiser clients as services. Technologies which report on campaign performance are drawing in data from the third party ad server and other data sources. Technologies for buying or trading, which connect to exchanges and Publisher networks, are built by Agencies for exclusive use by their Advertisers and are used to attract more Advertiser business. Technologies such as Agency-owned planning tools are integrated with third party ad servers to speed up the process of moving from planning to trafficking. Such tools add a competitive advantage to one Media Agency over another and are an important consideration for any Advertiser selecting a new Media Agency to work with.

Media Agency data analysts
The Media Agency roles and responsibilities in campaign setup and management from research and planning through to buying and trafficking help demonstrate why most Advertisers choose to employ the expertise of such stakeholders. The role of the Media Agency does not end when the campaign begins but goes on to add strength through expertise in optimization, reporting and the on-going trading of advertising space. It is important to state that all media agencies are different, even within the same holding group Agencies are known to follow different processes and workflows. Some Agencies have account teams but then split out the responsibility of campaign management by channel by having additional teams in place such as a search team or an affiliate team. Further to this some Agencies who have a particular focus on data-heavy strategies, employ a team of expert analysts; who build data processing systems and retrieve the most granular type of reporting from the ad-servers during and post campaign. Media Agency teams that have analytics resources are on the rise as the industry ventures towards greater use of what has been termed

"big data": making use of large datasets such as ad server cookie level reporting to make business decisions.

It is at the Advertiser's discretion to employ the Media Agency and by the same token, the Advertiser chooses a Creative Agency. Creative Agencies are also not essential to every digital campaign; campaigns which run on channels like Search do not require a designer or developer to produce an aesthetic visual as an advert. Those channels which do have this need such as Video or Standard display bring the Creative Agency into the campaign setup process at the point that the Media Agency issues the media plan. Since this document details the specifications that the creative Agency will need to build to, the Creative Agency cannot begin their work without it. The next section will go into detail about the role and workflows undertaken by the Creative Agency.

Chapter 2 - Section 3
Creative Shop

- The Creative Shop
    - Relationship with the Advertiser
    - Relationship with Publishers
    - Relationship with the Media Agency
- Marketing budget for Creative development
- Roles and Responsibilities
    - *Figure 2.7*
- Getting the Design right
- Development and Production
- Flash authoring
    - *Figure 2.7.1*
- Actionscript
    - *Figure 2.7.2*
- Components, building and publishing ads
- Ad storage systems (Asset Management Systems)
- Localisation
- Standard QA
    - *Figure 2.7.3*
- Uploading Creative
- What's in a name?
- Creative Inspiration

This section defines the role of the Creative Agency and examines the responsibilities that exist there to support the digital campaign setup process. In addition we will take a deeper look into the definition of an ad or creative to understand the skills, tools and systems that are mandatory for the Creative Agency to support the Advertiser, Media Agency and Publisher.

The Creative Shop

Visual ads are designed and created by a "creative team". Such a team can either sit inside a Media Agency or more commonly sits within an independent Agency of its own right; the "Creative Agency" (sometimes referred to as a "Creative Shop" in the US). When a media agency or advertiser,chooses to plan or setup a digital ad campaign that does not utilise complex visual elements (such as with campaigns setup in the search channel) the services of the creative Agency are not required. As the possible digital advertising channels have evolved, so has the role of the Creative Agency: as we will see later, channels like Dynamic Creative Optimization (DCO) have completely changed the perception and the expectation around Creative Agency involvement throughout the campaign.

Creative Agencies build their predominant relationships with the Advertisers who employee them. Relationships with Media Agencies and Publishers, which are just as crucial to maintain for successful ad campaigns exist for different purposes with separate dialogue required. Let's look at this balance more closely:

1. **Relationship with the Advertiser**
   To first work with an Advertiser, a Creative Agency will attempt to advertise the skills of their designers and developers by obtaining the Advertiser campaign brief and preparing mock-ups to showcase to the potential "client". Mock-ups will be the creative realisation of the brief into the story-board of the campaign in an aesthetically innovative way. This work is undertaken for free at cost to the Creative Agency that can hope to win the attention of the Advertiser and if successful; their business for the campaign and beyond. Such tactics make the creative Agency work very competitive and the protection and limited distribution of the campaign brief, crucial to the existing Creative Agency.

Independent Creative Agencies have very close relationships with the Advertiser, more so than they would have with the Media Agency in fact and this all comes down to the development of a brand. In brand development, the Creative Agency is employed to birth the Advertiser's brand standards and guidelines. By keeping a consistent loyalty to those guidelines the consumer experience of the Advertiser and the ads that they encounter enable far stronger brand and product recall. The Advertiser therefore invites the Creative Agency to approve and develop all creative in digital ad campaigns as well to adhere to the creed of the brand. Having developed such standards, the Advertiser will look to use the Creative Agency to approve work undertaken by the Media Agency that may impede on the perception of the brand. Alongside this auditing process the actual output delivered by the Creative Agency is referred to in the digital world as "the creative" and consists of images and content in the form of digital files.

2. **Relationship with Publishers**

Publishers own the infrastructure of the web, the sites and ad spaces themselves. The Publisher can restrict the type of ads that can be used on their websites. These limitations or "specifications" and mostly unique to each website and legally must be observed in order for the creative to "go live" on a Publisher website. The legal agreement to adhere to such specifications is laid out when the Media Agency commit to the purchase of the inventory.

Publishers will check the creative thoroughly before the campaign goes live. Creative developers that may try to side-step the specifications using clever coding will have their creative caught by the Publisher operations team. On the one hand the inclusion of any limitations to the Creative Agency may seem like a barrier to expression and creativity, but the Publisher must look out for its own interests. These interests

include the preservation of its own design and the specifications ensure that the creative is as visually effective as it can be in the publisher created environment of their webpage.

Creative Agencies and Publishers will often work hand in hand on expressive and elaborate executions that, for one campaign or two, are granted permission to breach Publisher specifications. Such exercises can mean good publicity for both parties, as well as impactful branding for the Advertiser.

3. **Relationship with the Media Agency**
   Commission an artist to paint and he will paint until the masterpiece is complete, his own interpretation and style at the helm. For the artist at the creative agency, a party that is not doing the commissioning could be viewed as a hindrance to the creative process. However, when a Creative Agency develops the "creative" without *some* collaboration from the Media Agency, things go wrong, crucial details are missed or the intent of the campaign goals could be misunderstood.

   The answer is not to involve both parties at every step in the creative development process. The relationship between Media Agency and Creative Agency is a balancing act, because the Creative Agency needs to retain creative directorship while still inviting feedback and contextual campaign information from the Media Agency. The right balance is best struck by working with both a Media and Creative Agency that have a strong *existing* relationship. These optimum relations give success to a campaign far beyond the use of the traditional Media Agency-Creative Agency engagement; the use of the Media Plan.

As we have seen, a media plan is an important tool to align the Media Agency" buys" with the creative that needs to be developed for them. Laid out in a template that might be foreign

Marketing budget for Creative development

Visual ads, animations and images are some of the most expensive components of a digital advertising campaign. The campaign budget is set aside early on for the development of creative to the specifications of the final media plan. In turn Creative Agencies can charge hourly rates or fixed costs reflecting the output of ads, either by tallying them up or by charging based on complexity or a combination of both. Campaigns which have more ads or more complex ad functions can therefore cost a lot more and hit the campaign budget harder before the campaign has ever begun.

Since campaigns can also change over time, the need can arise fo more assets and ads, so it is sensible to set aside part of the marketing budget for additional, unforeseen creative costs. In addition to this creating new ads for every new ad campaign can get very costly for the Advertiser. Over an Advertiser's lifetime, millions are spent on building ads from scratch.

Third party ad serving technology can assist in this area: the technical solution is to develop creative that utilises DCO (Dynamic Creative Optimization). We will cover this off in full detail in Chapter 3 but Creative Agencies that aim to develop DCO ads can provide a very significant saving in the long term for Advertisers looking to run numerous digital campaigns over the course of a year or multiple years. Millions can be reduced to hundreds of thousands of dollars without sacrificing creativity.

When the budget has been set aside, the Creative Agency teams can begin working on the creative for the campaign.

Roles and Responsibilities

A creative "build" requires the efforts of two distinctive roles at the Creative Agency. These roles are the designer and the

developer (or producer). The full Creative Agency's structure is formed as follows:

- The **Account Management Team**: This is the team that handles the relationship with the Advertiser directly and creates the briefs for the creative team. This team is in charge of managing expectations between the Advertiser (referred to as the client) and the creative team.

- The **Creative Team:** Formed of the Art Director and the Copywriter. The Art Director (designer) is in charge of producing the graphics and leads the designing of creative containing animation. The copywriter has the responsibility of coming up with the text (or "copy") in the ads. These roles collude to come up with the overall concept for the final creative output, "story-boarding" them for the production team to develop.

- The **Production Team:** The production team's structure varies from Agency to Agency. Some big Creative Agencies (like *Saatchi & Saatchi*) tend to outsource most of their production work to specialist "Production Shops". In general, a production team would have front-end Developers and back-end developers. Front-end developers would in charge of coding the "look-and-feel" part of the Advertiser's websites using Flash, JavaScript and HTML among other languages. The back-end Developers in charge of coding in server side languages for Advertiser websites or in complex in-ad languages like Actionscript - which we will cover off in more detail in the next few pages.

  The Developers would receive the "art" or "creative" from the creative team, have a discussion to agree on feasibility, to code and develop the finalised output which is usually a Flash file called a *.swf

In some instances, standard display campaigns are easy enough to be developed directly by the art directors in the creative team who generally have the required Flash animation skills to do the work. Here there is no need to employee the work of the production team, who might just be involved in more complex executions like Rich Media.

The creative Agency teams work together in the creation of creative for the campaign. At a high level the process is as follows (this process may differ from Agency to Agency):

1. The Account Management team solicits the campaign brief from the Advertiser
2. The Account management team receives the media plan from the Media Agency
3. The Account Management team supplies these documents to the Creative team
4. **The Creative team design the creative and build out the assets**
5. The Creative team supply the assets to the Production team
6. **The Production team author, develop and build the final functioning creative**
7. **The Production team QA the creative**
8. The Production team upload the creative to the third party ad server
9. The Production team supply the reference to the creative, to the Media Agency for trafficking

The Account Management team tend to have less engagement with the technology; it is the Creative and Production teams that undertake the production-line style processing bringing a digital display ad to life. The journey is one from Design to Development to QA (testing) and then to upload, the remainder of this section will explore the technologies that support this creative process.

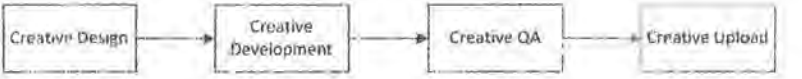


Figure 2.7: *The four crucial stages in the creation and purposing of the Creative Agency output.*

Getting the Design right

The creative team design and create the images and assets for use in the final piece of creative. Images are often created to high resolutions so that they can be used for multiple ad dimensions. Designers achieve this feat using what is referred to as vector-based design. Vector based design allows images to be scaled down or up without a loss of image quality. Designers will use programs like Adobe Photoshop to produce these assets but they are not free to develop very high quality images for all ads. As we touched on earlier, this is due to Publisher specifications and almost all Publishers across the web universally agree that standard display ads should be no more than around 40-50kb in file size. The creative designer needs to be very careful tweaking images and assets to ensure that the creative remains aesthetically pleasing while still meeting the requirements on file size.

Rich media offers designers a more liberal opportunity to include much higher resolution imagery as the constraints on file size are lifted (although these too will be subject to some upper limitations). Video creation also falls into the remit of the designer and the combinations of images, video and even high resolution imagery are the resulting creations that the designer will pass along for development or animation to the production team.

The final output of the creative development process will be one of the following file outcomes depending on what the media plan requires:

Standard Display Campaign - A minimum of one *.gif file per ad required. A maximum of one *.gif file and one *.swf file per ad required.
Rich Media Campaign - A minimum of one *.gif file and one*.swf file per ad required. A maximum of multiple files of many different types including video files.

In Chapter 3; Campaign Setup by Channel, we will look at why these precise configurations are required. For the this section we will focus on exactly what *.gif files are and what *.swf files are and how they are created.

Gif's as most readers might be aware are just ordinary images. They are used as a file type because they have a legacy for having good compression rate and pack nicely into a small file for the Creative Agency to work with and deliver. The gifs would be built by the creative team in a program such as *Photoshop*. The creative team will work on other images and drawings as well but these are not outputted to a final file. Instead they make up part of the production process known as creative authoring. It is the full process of authoring which outputs the other file type mentioned: the *.swf file.

Development and Production
The production team work with the functionality of the creative images that have been passed along by the creative team. Creative for ad campaigns is rarely a static (still) image; it has lots of moving parts, it might be animated or contain controls for interactivity. It may respond to external stimulus or information, and behave accordingly under certain conditions. The production team produce the final files, typically working inside Adobe's Flash

authoring program so that the ad contains this functionality. So what is flash authoring?

Flash authoring

As we touched on briefly in Chapter 1, *Adobe Flash* is a program that allows the creation of dynamic multimedia; typically outputting an animated file that can be played in any browser where the flash player "plugin" has been installed. With vast dominance on almost all web accessing machines (although less so in the now growing tablet and mobile spaces): Flash has been the development environment of choice for digital advertising.

The production team developers will have copy of Flash installed on their local machines and will spend many hours using it to create the desired functional effects inside the creative. The area designated for the assembly of the creative is referred to as the "stage", which displays the current "frame" that is being worked on. Frames can be layered on top of one another or become hidden and as a default the Flash animation would contain a series of frames representing the progression of the animation. Just as with an image created using *Photoshop*, the program has a selection of artistic tools, visual filters and effects. Among the most useful of these effects are "tweens" (pre-coded animation paths to change shapes or provide motion) and the availability of Actionscript.

The authoring tool also utilises a special storage area of imported images, shapes and code called the "library" which can be unique to each animation. Items in the library used in the animation are formally called "assets". Today these are the atoms of the ad serving and building world, there is no entity smaller to play with, move and transport than "assets" but today these are mostly trapped into the authoring tool and cannot be released from the cage of the resulting creative. In Chapter 3 - Dynamic Creative, we

will look at how new forms of creative development are changing this.

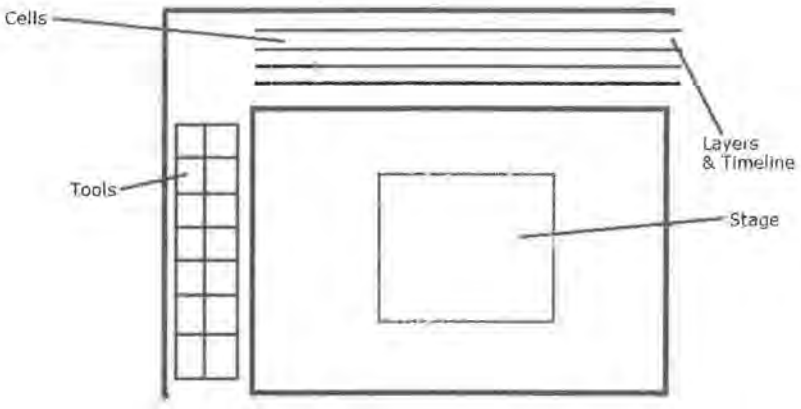


Figure 2.7.1: *A mock-up of the Flash authoring tool, the creative developer's workshop. Many hours will be spent inside this interface to compile the creative in an *.FLA file before it is published inside the same tool to a *.SWF file. The tool is very flexible allowing for the use of designer-ready tools as well as Actionscript code, which can be trigger by anything from the passage of time to the clicking of a button.*

Actionscript

Actionscript is a cousin of JavaScript, much of the code looks similar and with a good grasp of JavaScript, Actionscript is not hard to pick up. It is a very versatile language that can be used to mathematically produce images and animation rather than importing creative assets. The advantage of this method is that the overall file weight of the final creative will be reduced (to leave room for creative flare with more code or more assets).

The creative flare that Actionscript provides the developer with should be met with the objective to offer a user experience that meets what we call in the industry: "the lowest common

denominator". Stolen from the pages of a high school maths textbook, Lowest Common Denominator means that creative should be developed to allow even users with the slowest computers and internet connections to have a great experience. The Publisher specifications are designed to ensure this so when building to those specifications the developer compares the specifications across all of the sites on the media plan and produces creative that adheres to the strictest guidelines. For instance; if the strictest specifications demanded a 2mb file but other publishers were more liberal, allowing file sizes up to 5mb then the developer would build the creative to a 2mb specification so that the same file could be used across all of the publishers. This is far more efficient on the creative side that creating a version for each Publisher to corresponding specifications.

On top of the Publishers demands, the developer is recommended to consider guidelines that all Publishers have missed from their specifications but do affect the user's experience with the ad. One such example is the reliance on Actionscript within the creative to "draw" objects and shapes in the creative in real-time when the user encounters the ad. Although this method reduces file weight, it can be a hefty drain on the RAM of the user's machine and slow down the performance of their browser. Such ads look outstanding on high specification computers, but as a best practice the Lowest Common Denominator of the user's own device is an important consideration.

So far we have described the design and build of the creative as a freehand operation, undertaken by the developer inside the Flash authoring tool. Third party ad servers have always given this freedom to flash developers; allowed them to use raw code and start the build process for the creative from scratch for each project, providing help files, training and manuals to assist the build. As the industry has evolved, the skills required to work with

code in this manner have become ever more demanding and complex so the number of developers in the industry has shrunk while the number of designers has increased. As we have seen this has resulted in the out-sourcing of the development work to other Agencies. Another growing trend is to outsource to freelance developers.

To assist less-skilled or busier developers, the third party ad server provides assistance inside the Flash authoring tool itself by supplying pre-built templates and using small modules called "Components" to give the Flash authoring tool a direct connection to the third party ad server. This applies to all the processes we have looked at already: Design and Production, but also extends to QA and facilitates; an easy upload into the third party ad server User interface.



Figure 2.7.2: *Components open as a separate panel in the Flash Authoring tool and allow for the fast creation of creative to ad server specifications quickly and at scale for campaigns where creative needs to be produced in bulk. Creative producers that are already familiar with components like DG/Mediamind's workshop tool can be expected to deliver creative faster than those creative and developers that are not.*

Components, building and publishing ads

Third party ad serving providers offer a range of components to make it easier for developers to build ads. Components come in packages which are usually free to install and use. The *DG Mediamind* package is called "the workshop" and installs onto the developers machine. Updates are usually supplied every few months just like any other computer program to ensure the best creative kit is available to the developer. Among the components will be a pre-built video player as well as a pre-built data capture form. Perhaps the most useful among the components is a library of pre-designed templates (something that *DG Mediamind* refers to as "Blocks"). Loading one of these templates will unpack onto the authoring stage a near-complete Flash project. Clicking a few settings and making some adjustments allows the developer to tweak the template to match the Publisher specifications for the creative. The developer will then drag and drop the assets or images created by the design team into the same project and assemble them so that the frames of the final creative can be played through in running order.

The described project that the developer is working on is actually a file called an *.FLA file. An FLA file cannot be used outside of the Flash authoring tool but stores the designers work in the design stage and when the file is passed along to the production team, it stores their work too. When the creative has been finished, the developer can choose to "Publish" the creative which allows the Flash Authoring tool to save the file to the developer's machine as a *.SWF file. The *.SWF files are the creative output that are sent on to the next stage in the setup of the ad campaign. *.FLA files should not leave the creative developer's machine but they should be retained in case problems arise later on (more information on this can be found in Chapter 3- Rich media). *.SWF files cannot be adjusted later on without opening up the *.FLA file and choosing the "publish" option in the Flash authoring interface. FLA files are useful to keep hold of. Should a *.SWF ever

be misplaced by any party, the FLA can be used to generate another copy of the *.SWF file.

Sometimes an Advertiser will run campaigns quite scarcely and to lower costs, will use and adapt old creative for future campaigns. For flexibility to use other developers and Creative Agencies the Advertiser should ensure that all files used in the creation of the creative are supplied to the Advertiser directly at the end of the development process. This includes all FLA's, SWF's, Images, Videos and Fonts. (Font's, which facilitate the design of written content in a creative actually only display that way on the developers machine because they are loaded into the FLA file separately). When the production team converts the files from FLA to SWF the fonts become embedded in the SWF so that they display correctly. If a production team were to out-source development work to a freelancer who supplied all files except the Fonts at the end of the project, were the production team wish to make adjustment to the creative by "re-publishing" the SWF from the FLA file, any text used for an uncommon font would display incorrectly. The freelancer would have to be sure to supply the font file he had used at the end of the project.

Ad storage systems (Asset Management Systems)
High performing creative production teams are those that are most organised when so many files come into play. If each creative has three files for instance, and ten different dimensions are created, in ten different languages, there are 300 files being designed and built. This doesn't take into account the more raw assets like original video files that the design team may have been working on. Clearly there is a need for intelligent file storage systems and in the world of ad development these are called Asset Management Systems.

Some parties unfamiliar with the inner workings of an ad server may be unfamiliar with the storage facilitates that it has available. The third party ad server is not designed as a storage facility for

creative. Creative uploaded to the ad server cannot easily be downloaded again, and the ad server does not support the upload of file types used for development such as FLA files or PSD files (the files used by the designer in the creation of layered images in Photoshop). There is however an area called the "creative library". This is tree of folders inside the third party ad serving interface that can visualise and arrange creative before it is picked up for the trafficking process (more on trafficking in Chapter 3 - Display ad serving). The creative library can be seen as 'the drop off point' for the creative, but it is an area relatively rudimentary compared to true Asset management systems.

Asset management systems are not designed for mass ad *distribution* like the third party ad server but provide secure archiving so that historical ads can be retrieved later. These systems also provide the ability to tag or label creative in intuitive ways (which is called the application of "meta data") so that they can be easily retrieved, recovered and assembled later on. Assessment Management systems are also essential to Creative Agencies that work across many borders in many different countries.

Localisation

Big advertisers with campaigns that run in multiple countries and languages will employ one Creative Agency to produce a selection of core templates at a central point globally which is then distributed to Creative Agencies in local countries where the ad campaign will be running. This process is called "Localisation". The Assessment Management system allows for controlled storage and distribution to these local teams where the local creative designers will adjust the creative both for language and to local Publisher specifications. This way nuisances in the media plan at the local country level are not completely restricted to stringent branding guidelines but have some flexibility. The Asset Management system allows the globalised Advertiser to retain a degree of control by feeding brand guidelines through to local

teams and approving localised creative to minimum brand guidelines. Global campaigns co-ordinated in this way often follow a long chain of brand guideline approval from the original global marketing team to ensure that even for the smallest Publishers; creative and brand guidelines are adhered to. Asset Management Systems are still in their infancy but help the production team to keep their desks tidy when the number of files in development escalates.

Standard QA

Once the production team have finished developing and "publishing" the *.SWF they must test it to make sure it works accordingly. This process is called "QA" (Quality Assurance) and it can be a complex undertaking if a lot of files are involved and the Advertiser wants to be very thorough about the final creative delivery. To be thorough there would be four different QA processes that would occur in the lifetime of an ad:

1. Local Machine QA by the creative producer
2. QA in the third party ad server ad preview interface
3. QA by the Publisher (to check it against the Publisher specifications)

The reason to do so many checks is that the *.SWF is being transported and embedded in an environment where there are a lot of other things going on (the Publisher website). Particularly when developing Rich Media ads, more can go wrong so a longer and a more stringent QA process is recommended. When it comes to the creation of Standard Display *.SWFs the development cycle will be much shorter and the ads far less complex. In instances where there are thousands of ads to check, step (2.) might be missed out entirely or the work out-sourced or undertaken by spot-checking the files. Typically a good creative production team will undertake the first QA check on this list.

When performing a QA the production team are checking that the creative is functioning and display as expected. On a local machine this can easily be achieved by publishing the *.swf file and viewing it in the browser window or in the authoring tool preview screen. A true QA involved checks that all the interactive functionality is working, that the ad has no excess loops or glitches and that it meets the lowest common denominator specifications for the Publisher. The most crucial check is the functionality of the ability to click, and the developer will encode this into the Actionscript of the creative. This specific script is called "the clickTag" and the third party ad server components ensure that the clickTag is always embedded in the creative from the outset of the production process behind the scenes. Despite this, the absence or incorrect coding of the clickTag is the most common reason that ads get rejected in the QA process in (2.) and this is because the clickTag code is still not standardised among the ad servers. It is a little different depending on which one gets used.

```
on (release) {getURL(clicktag, "_blank");}
```

Figure 2.7.3: *An example of a clickTag coded in Actionscript 2.0. Ad servers are very particular about the precise usage of their specific clickTag, creative which fails to include the correct clickTag code can be rejected by the ad server. This failure can cause significant delays to the strict timelines of a digital ad campaign.*

Uploading the creative

With the creative QA'd by the production team it can now be uploaded to the third party ad server. Once uploaded the creative will live in the creative library until the Media Agency assigned "trafficker" selects it for use in the setup of the ad campaign. The third party ad-server supplied components have built-in the functionality which allows the creative to upload the final, tested

creative to the ad server. To do this the developer needs to have a username and password setup which will ensure the creative is uploaded to a designated private creative library. The third party ad server will issue these details to the Creative Agency with the approval of the Advertiser or Media Agency. The final creative can be uploaded as single files or in bulk. Some ad servers allow creative in several files to be zipped up in an archive file or zip file and uploaded to the third party ad server that way. The upload process causes the third party ad server to solicit a reference number. This is called an "Ad ID" or "Creative ID" and the component which facilities the upload of the creative displays this ID to the developer for their reference. The creative developer will then supply this number to the Media Agency trafficker. Some ad servers offer additional QA services, just to be sure that the creative at least meets ad server specifications (which will mean that the creative contains the specified clickTag and that it contains all the required methods that any creative built using the ad server supplied components will contain in any case).

What's in a name?

In the process of the uploading the *.SWF file, the developer should apply the same naming convention that will be used when naming and uploading all the other *.SWF files to the ad server. The best naming conventions are reached in collaboration between the developer and the Media Agency because the creative file name will eventually find its way into ad server reporting, and ad performance can be analysed by the creative file name. If a *.SWF file name is set up to contain information like "creative message", then the performance of creative using a specific message can be determined later on for example.

The role of production team ends with the uploading of the creative. The assistance of this team may be called upon at a later date to undertake the same tasks and produce further *.SWF files or adapts the existing creative by republishing.

The Creative Agency are occasionally asked to participate further in the setup of the campaign in the ad server interface itself but the role to play here is a negotiation between the role of the Media Agency trafficker and the creative control held by the creative producer. With the advent of DCO (see Chapter 3 - DCO) creative functionality and design can be controlled outside of the Flash authoring environment and in the ad server UI instead. It is the choice of the creative producer to take on more of a trafficking role to control the birth of new versions of the same creative.

Creative Inspiration

The canvas of the ad may not sound like the most glamorous of locations for the birth of a masterpiece but incredible creations are born every day the world over and these gems are increasingly recognised, adopted and copied. Such executions can win awards leading to fame for designers and production teams. The marriage of cutting edge technology and compelling artwork makes for some truly inspiring ads that put campaigns in the history books and make Agencies and Advertisers think twice about employing the most formidable creative talent. This is the reason to make a wise consideration for a Creative Agency and to harness a long term relationship with talented designers and developers.

Now that we have looked at the end to end process of creative design and production it is time to turn attention to the role of the Publishers, they are the last stakeholder we will focus on. Powerful and extremely different in the way they approach and utilise the web, they are the ally that every Advertiser quickly realises it needs to have. After all, without the ad space, there is nowhere to display the art produced by the Creative Agency and no audience to show it to.

Chapter 2 - Section 4
Publishers

- What are Publishers?
    - *Figure 2.8.1*
- How do Publishers sell their ad space?
- CPM
- CPC
- CPA
- Publisher Networks
    - *Figure 2.8.2*
    - Niche Networks
    - Blind Networks
- Publisher ad servers
    - *Figure 2.8.3*
    - Inventory
    - Publisher Specifications
    - Publisher QA
    - Publisher Ad Ops
        - Figure 2.8.4
        - Ad prioritisation
        - Frequency capping
        - Over and Under Delivery
        - Browser Targeting
- *Figure 2.8.5*
- How does Publisher reporting differ from third party reports?
- Discrepancies

What are Publishers?

The online advertising industry refers to the majority of the websites on the world wide web as Publisher websites. These are the websites that attract internet users; they are hubs for both audiences and content, often pumping vast sums of money into

creating such content, to attract audiences. The audience has an experience on a Publisher website of mostly free, rich information which brings the internet alive; the illusion is that the internet somehow offers a world of information for next to zero cost. The underlying truth of course is that content of the Publisher sites costs money to create, to store, to manage and it is through the eyeballs of those audience members which generate the true value. The audience attention rests on the content, but within peripheral vision sits digital advertising where Advertisers spend vast sums of money in the hope that ads will be seen, clicked and potentially used as a conduit to make a purchase.

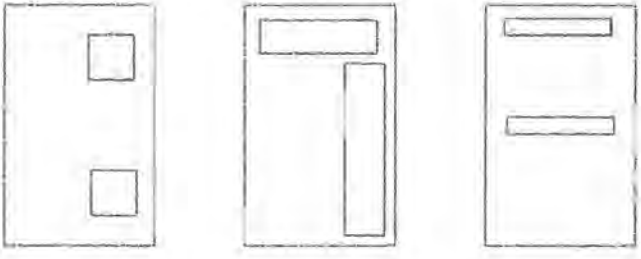


Figure 2.8.1 *: Example Publisher Pages with empty ad spots in varying sizes.*

Figure 2.8.1 shows an example of three webpages making up a very small fictional Publisher site. As a Publisher if I want to monetise my website without charging users to access these three pages then I have a total of 6 ad spots I can sell. If my Publisher site only has one user visit each page once over a period of a month then it's safe to say that I'm not going to make any money but if my Publisher site is very popular with millions of users visiting these three pages in a month then suddenly I can sell the ad space and make money. How much ad space do I have available to sell? The answer is lots; not just 6 ad spaces, but every time the page is loaded I could show the user a new ad. If each user sees just one page and there are a million users a month then my "inventory" is a million impressions.

The catch to this is that I actually have to predict my how many users or how much traffic I will have in the future to know how much ad space I have to sell and if I want to chop up this inventory and sell it to multiple Advertisers, how can I do it? These questions are solved by Publisher-side ad serving technology. Selling off the inventory impression by impression in this way is just one of the methods Publishers utilise to sell their ad space.

How do Publishers sell their ad space?

The Publisher sells the ad space as "inventory" on to Media Agency buyers (see Chapter 2 - Agency Planning) across three core "pricing models". The models are:

1. CPM or cost per 'mille', offering a thousand ad views or "impressions" at a price. CPM reflects the value of the audience seeing the ads; the further down the purchase funnel they are, the higher value the ad space; if the ad spot is seen more often, by a less targeted audience, the lower the CPM. Also the more valuable the Publisher is perceived to be or the more niche the audience; the higher the cost to buy the space. How does a Publisher decide on the CPM pricing? This is more complex than it may sound. Prices for ad inventory vary enormously between sites, pages and Publishers. Considerations for coming to a reasonable and fair price include any variable that might influence the user's visual exposure to an ad on the final purchase with that Advertiser. This will include: the Format of the ad, the audience type that it is shown to, the total amount of traffic that the ad spot is expected to be exposed to, the total time of exposure, the strength of the content and whether the Advertiser is looking to sell directly through the ad or just conduct an exercise in branding. Beyond these considerations the pricing needs to remain competitive with other publishers. In some cases Advertiser and Agencies help Publishers to determine the final

price based on the pricing of similar ad spots in the market in the process of negotiating the cost of the inventory. CPMs are not an exact science or uniform across the web. For Agencies and Advertisers it pays to get a quote for ad spots with several different Publishers to get a good guide on a market price.

2. CPC models dominated the early internet; forcing the Advertiser to pay for clicks not the number of times the ad is seen. Clicks are more valuable because the user demonstrates some intent to engage with the ad, product or brand. But clicks do not work in favour of the Publisher; if the creative is badly designed the Publisher takes the hit by showing an ad thousands of times with no clicks. As we will see later in this section the Publisher can also track the ad's click by calculating the likelihood for a standard user to click on the ad over the course of a campaign the Publisher can determine what a good price would be to charge per click. If a Publisher can get away with serving just a few ads but expects a click for almost every user, it will be more favourable to charge on a CPC. Where CPC models exist "long tail" Publishers have been known to try and dupe the ad serving counters employing small armies of internet robot programs to mimic users clicking. Such techniques do not take into account that ad servers have filtering layers in place that do not pass fraudulent click activity onto the reporting part of the ad server.

3. The last model is CPA or cost per acquisition. The presence of a conversion tracker on the Advertiser site (supplied by either the Publisher themselves or the third party ad server) ties a conversion back to the ad and publisher which drove that conversion, this can also be extended to a lead or sign up (more details on the pathway to conversion are covered in Chapter 4). For this model a "conversion" just means triggering a conversion tag" so this can mean the completion of any specific action. Due to the flexibility of this definition

there are pricing models that exist like "CPL" (Cost Per lead) where the action of generating the lead will fire a "conversion tag" - so in this sense CPA and CPL are the same.

Running a business selling ad space is never quite as simple as three pricing models. Others which I have encountered: Fixed Price, CPM combined with a frequency cap, FOC (free of charge), CPMV (cost per visual impression), CPE (Cost per Engagement), Rev-share models (Revenue-Share), Sponsorship (a model particularly popular in the Chinese market) and a whole range of deals where online ad space is packaged with ad space for different channels such as TV. New technologies in ad serving will mean lots more options for inventory purchase in the future.

Publisher Networks

With these models in hand the Publisher can approach their the media Agency buyer or Advertiser. Publishers must monetise their websites to offset the costs of producing and maintaining content and without subscription this can only be achieved through the sale of ad inventory. The Publisher typically has a sales team that packages up the inventory based on common knowledge or, on occasion, assumptions about the audience viewing it. For instance the Autos Section of a website is typically classed as inventory suitable for middle aged men in the market to purchase a vehicle. The Publisher sales team might choose to contact Autos Advertisers (or Media Agencies representing Autos Advertisers directly) to sell, closing the deal by having the new customer sign an IO (Insertion Order).

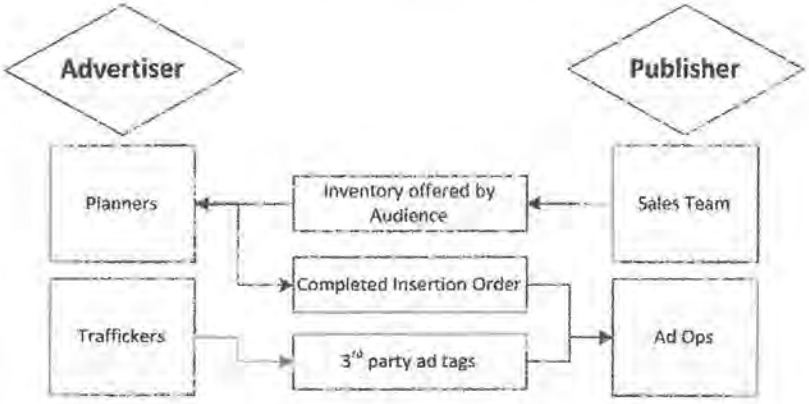


Figure 2.8.2 *: Campaign Setup Workflow between Advertiser or Media Agency and Publisher*

The completed IO is processed by a Publisher ad operations team who use a "Publisher ad server" to load ads and ad tags onto the Publisher website. A Publisher network acts as a sales and operations house for multiple Publishers, and normally only take on Publishers with a certain amount of available inventory and audiences which are in demand by Advertisers. For a Publisher, signing up to become part of a network means the security of outsourcing most of the advertising campaign operations, and in some instances contractual agreements are made to guarantee revenue.

Publisher networks hugely reduce wastage of inventory by being able to select from a greater pool of Advertisers and capping the user's exposure to a particular campaign or Advertisers messages across multiple sites. Networks having been going through big changes in the last two years, The biggest change of all has been the advent of Advertising Exchanges. Advertising Exchanges and SSPs are covered in more detail in Chapter 8 - Programmatic.

Niche Networks

A Network is a term thrown around a lot in the Publisher space. Mobile Network and Video Network are examples of Publisher Networks that are more channel specific which are becoming much more common in online advertising. Another type of network is the Affiliate Network and its complexity gives it a channel all of its own. To understand the difference briefly; Affiliate networks operate under more dynamic cost models than standard Publisher networks and offer a differentiated reach of the online media audience. As affiliation is actually considered a separate channel to both Search and Display; this book dedicates a complete section to it in Chapter 3 - Affiliates.

Blind Networks

Networks sell the available inventory but aren't always completely transparent about exactly which pages the ads will appear on or even which sites in their Publisher network they will be served on. With a blind network buy there is no way for the Advertiser to go to a page to see their ad when the campaign does go live. Only reporting will reveal the delivery statistics which is why using a third party ad server for these "buys" is essential.

The relationship between the Publisher and media buyer is facilitated through the agreement of an IO and since several Publishers will be on the Media plan for any given campaign, the job of determining exactly how much to spend with who can be determined by the relationship that the Publisher sales contact has with that Media Agency buyer. Once the relationship has been established it's time to utilise the technology to assist with the completion of the buying process.

Publisher ad servers

There are two sorts of Ad server. Firstly is the third party ad server (also known as the "Advertiser ad server" which we have been referencing to date) and secondly is "the Publisher ad server". The two are commonly confused but very few companies produce both technologies under the same branding; the exception to the rule is *DoubleClick*, offering *DFA (DART for Advertisers)* as the Advertiser ad server and *DFP (DART for Publishers)* as the Publisher ad server (and merging the two under a rebranding exercise which is taking *Google* some years to achieve).

A Publisher ad server is a platform that measures and manages the ad spaces on a Publisher web site, sites or domains. It is typically rolled out only to single sites or networks as opposed to Advertiser ad servers that have a much larger cross-Network reach. Publisher ad servers are most commonly accessed by the advertising operations team at the Network or Publisher.

Initially the Publisher ad server can be used to establish how many ad spots on the site are available to receive advertising currently and in the future. Sales houses and reps for Networks and Publishers, utilise the technology to package the available ad inventory for sale. Examples of Publisher Ad Servers include *Adtech*, *DFA* and *OpenX*.

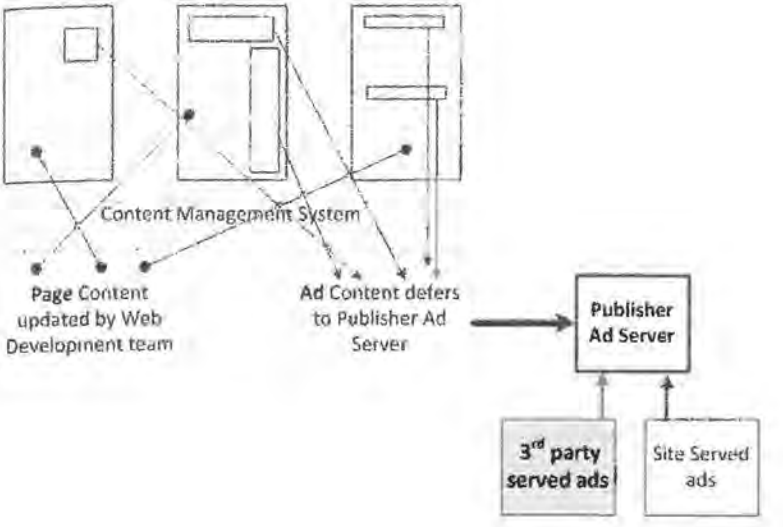


Figure 2.8.3 *: Management of Publisher Sites - Content Vs. Ads. Ads are then served by the Publisher ad server (such as DFA) being served themselves by either a locally served ad or a third party served ad.*

Inventory

Once available inventory has been sold to an Advertiser or Agency; the advertising operations team on the Publisher side awaits the details of the signed IO, including the ad tags "trafficked" from the Advertiser ad server. Sometimes ad tags are not supplied, such as in instances where Publishers have sold inventory and agreed to serve the assets directly from their Publisher ad server or domain. This is also referred to as "site serving" and, in some instances offers a lower overhead to the Advertiser while compromising on the benefits of de-duplication and the third party report verification received from the third party ad server. In 2014 in the UK, some traditional TV broadcasters continue to only site-serve "instream" video ads rather than receive third party tags (more on Instream in Chapter 3). Ads which eventually appear on Publisher pages must be built

to explicit specifications so as not to inhibit the user experience on the Publisher site.

Publisher Specifications

Publishers will publicly release specifications according to which ads need to be built by the Creative Agency. Public availability ensures that there can be no excuse for a Creative Agency with the media plan in hand to develop an ad which will not be rejected by the Publisher before the campaign goes live. Available inventory is also sold under the strict conditions that the Creative Agency adheres to the published specifications. For Standard Display ads the specifications usually conform to the industry accepted specifications for Display ads, set out by the IAB (http://www.iab.net/guidelines/508676/508767/displayguidelines). The IAB advise that Publishers place restrictions on file size (40-50kb is a standard file size for a standard Display ad) as well as the number of loops that the ad is allowed to have and limitations on "automatic play" and audio volume. Specifications also clearly provide creative guidelines for the content of the ad, for its suitability on the Publisher site, as well as rules for ad dimensions (please see Chapter 3 - Display Ads and Rich Media ads for more details).

Publisher QA

At the start of the campaign setup the Publisher ad operations team will receive the ads from the Media Agency or Advertiser (either as third party ad tags or raw creative). The first step in the setup process is to undertake a period of testing called the Publisher QA process. This will involve loading the ads into a test environment which will replicate as closely as possible the onsite experience of the loading of the actual ad. This test environment, also called a "staging area" looks and acts very much like the live Publisher site but allows the Publisher to:

a. Check the ad conforms to the Publisher specifications by checking page load speed or performance, looping and sizing
b. Checking that the ad does not contain any malicious code (including malware)
c. Checking that the ad does not have a bad interaction with other elements on the page (failures of the ad to perhaps load onto the correct "layer" of the page can be over-ridden in the Publisher's CMS system).

Publisher Ad Ops

The advertising operations team load the ad tags or assets into the Publisher ad server (providing they have passed QA) and apply various rules to their delivery such the agreed geographic location of the users IP address, ad prioritisation and frequency capping. Specific ad delivery ensures the ads appear where the Advertiser has paid for them to appear. For some agreements this may not be a specific page but rather "run of network" or RON which infers a randomised delivery to any ad spots (of the matching ad dimensions) anywhere on the network of sites connected to the Publisher ad server account. This same collection of settings gives the operations team the chance to ensure that the ad is delivered only to specific browser types or to specific IP addresses (or geo locations).



Figure 2.8.4 *: Publisher/Network Ad Operations. On the right of the diagram the Publisher team (sales and operations) liaise with the media buyer to share the available inventory and complete the purchase in advance of a particular date. This same team can access the Publisher ad server interface from their machines to determine what is available and to control what gets delivered and when.*

Ad prioritisation

A Publisher has the power to prioritise which ads to show to users based on how high a priority those ads are compared to other ads on the network. Prioritisation is used to over-ride the delivery of ads which are worth more money to the Publisher over a shorter period of time or to show ads of a lower worth when inventory is unsold. Such ads are referred to as "house ads" and may not be ads at all but visual assets used to direct users to particular parts of the Publisher website.

Frequency capping

Frequency capping is a more powerful feature in the Publisher ad serving world than for the third party ad server. Set the frequency cap in the latter, and backup images will be displayed when the cap is reached (as the third party ad

server always needs to serve something). But set the frequency cap in the Publisher ad server and upon hitting the cap, the user will be shown an ad for a completely different Advertiser; therefore more effectively capping the user exposure to a single Advertiser. Frequency caps are usually set to limit the user's exposure to just a campaign to a handful of "impressions" over a given time period of twenty four hours or less. Publishers are able to use such settings to their advantage to push more ads for a particular Advertiser over a shorter period of time. It is important therefore for the Advertiser to specify to the Publisher when booking their campaign the requirement of an effective frequency cap. Here the third party ad server plays a strong role by allowing the Advertiser to report on the effective number of times a user should see an ad in a given time period to ensure they convert, click or perform a desired function. Not enforcing a frequency cap can see the agreed purchase of inventory "burned through" in a very short period of time. Also setting the cap too long may impact the effective reach (users may not see an ad often enough) so it is an essential balancing act to set the cap correctly.

Over and Under Delivery

A Publisher ad server is supposed to be an effective tool in making sure that the number of ad spots sold to the Advertiser matches the number of ads displayed to users during the agreed time frame. Due to the complexity of many different ad campaigns (inventory lookups, page dynamics and sudden site traffic spikes as well as ad load discrepancies) a campaign can under-deliver (which is a burden to the Advertiser who is looking to display all the ads they paid for) or over-deliver (a burden on the Publisher which may see their agreed margins on sold inventory, diminish). Over and Under delivery can have knock-on effects to future booked inventory by cannibalising available space or leaving valuable gaps where there were previously

booked campaigns. Contractual compromises between Publishers and Advertisers are common place and now an accepted, but still a frustrating occurrence in the planning and buying of campaigns.

Browser Targeting
Publisher ad servers have the ability to target users based on their browser and IP address. The most common use for this functionality is to ensure that the ads are being delivered to IP addresses in the country where "audience" sold to the buyer is described as being located. This ensures that users that arrive to the Publisher site from somewhere else in the world are not shown ads by Advertisers that may have little or no presence in that country.

Once the Publisher ad operations team has processed the IO and inputted the campaign details into the Publisher Ad server interface, the ad campaign will go live at the set time and the Publisher Ad Server will work to deliver the ads evenly over the given time period (or however the delivery has been setup). At the last stage in the setup of each ad, the ad operations team will include small pieces of code called "macros" which allow the Publisher ad server to track clicks and prevent ads from caching on the local or remote machines as the ads are delivered. [As a side note: caching is a concept whereby a computer will store a copy of an image to save time and bandwidth by having it load from local memory rather that pulling the information down from the server for each load. It is important to prevent this from happening with ad serving because if the trafficker makes changes to the setup of the campaign or the ad itself, the change will not appear everywhere because the historical image will still be displayed]. The mechanism used to prevent caching is a concept that exists in the third party ad server as well as the Publisher ad server. Recording clicks is also something duplicated in both systems. When the ad operations team is finished with the trafficker's ad tag it can end up looking like this (nuances of this

particular method of adding macros differs among the various Publisher ad servers so this example is not accurate for all Publisher facing systems):

Before

```
<a href="http://altfarm.mediaplex.com/ad/ck/1234-12345-1234-1?mpt=[CACHEBUSTER]">
<img src="http://altfarm.mediaplex.com/ad/bn/1234-12345-1234-1?mpt=[CACHEBUSTER]" alt="Click Here"></a>
```

After

```
<a href="%chttp%3a%2f%2faltfarm.mediaplex.com/ad/ck/1234-12345-1234-1?mpt%n">
<img src="http://altfarm.mediaplex.com/ad/bn/1234-12345-1234-1?mpt=%n" alt="Click Here"></a>
```

Figure 2.8.5 *: In this example the ad ops team receives the ad tag from the Media Agency trafficker (the trafficker has sent this tag from the third party ad server Mediaplex onto the Publisher in question) The ad ops team amends the tag with a cache buster (%n) and a click tracking macro (%c) other character amendments are made to make the tag suitable for serving from the Publisher Ad server DFA*

To ensure that clicks are tracked accurately, some Publisher sites and networks include instructions in the Publisher specifications that the clickTag is coded in a particular way. We have touched the issue of clickTag inclusion the previous section in this chapter (The Creative Agency) but it is possible that the Publisher requirements regarding the use of the clickTag and the ad server requirements conflict so that it is not possible to fulfil both requirements. In this instance the Media Agency must inform the Creative Agency as to which set of specifications to build to, essentially allowing the reporting of clicks on the ad server type (the third party ad server) to overall the counting of clicks in the

Publisher ad server. When in doubt it is always recommend to build to the ad server specifications on the use of the clickTag. As the campaign gets underway, the ad servers (of both types) begin to store data and generate reporting.

How does Publisher reporting differ from third party reports?
As with the third party ad server, at the end of a campaign or during, the Publisher ad server can deliver a report to show how the ads performed and what was delivered to the user. The actual technical mechanism that counts these statistics actually counts them at slightly different times between these different ad server types (differing anywhere between one hundred and several thousand microseconds). This means that if you were to collect a delivery report from the third party ad server and compare it with the one from the Publisher ad server you are likely to see a discrepancy between the numbers. It is the role of the Publisher Ad Operations team to supply the reporting to the Media Agency (who may or may not use as a reference to optimize their campaign). The Ad Operations team can also pull the reports from the Publisher ad serving interface to check that the campaign is delivering correctly and make manual adjustments to the delivery settings to bring the actual delivery closer to the numbers promised in the IO.

Discrepancies
The IAB has created standards around acceptable limits of discrepancy (between the figures reported between different counting systems) which stand at five to ten per cent (plus or minus). Larger discrepancies can either be investigated, or the Advertiser can contractually agree with the Publisher to always pay out on the third party ad serving figures where the counting methodology of the third party ad server has been verified by the IAB. Discrepancy investigations by the support team at the third party ad server alongside the Publisher can take a long time, as the data sets need to be as comparable as possible and the

technologies may need to be tested side by side over time to uncover potential issues. The most common cause of discrepancies between Publisher and Advertiser technology is human error, where tags are implemented incorrectly or reports are pulled against the wrong tags, dates or placements, all of which are easily remedied by tighter processes and better communication between all parties involved in the campaign. Many third party ad servers will allow Publisher to have their own access to third party reporting directly so that they can keep an eye on delivery that way. Publishers accessing third party ad server reporting systems will be able to see all of the campaigns running on their network or site which may include multiple Advertisers, campaigns and Media Agencies.

Providing the ad campaign goes off without a hitch, the ads get delivered to the Publisher's audience over the agreed time frame and under the agreed conditions, the Media Agency is a step further towards meeting the advertiser's campaign goals. We have begun looking into the setup of a digital advertising campaign by examining each of the four major stakeholders and their connection to each other. The underpinning technical detail that achieves the feat of digital ad campaign delivery requires a much deeper understanding on technical detail, flushing out the work that the trafficker must undertake. Each channel is setup differently but for a third party ad server to be used to its full potential it must be used to track all channels that make up an Advertisers marketing mix.

Part B

Campaign Setup by Channel

Part B
Campaign Setup by Channel

- *Figure 2.8.6*
- How does the end to end Campaign Management workflow work?
- Billing, Electronic IO and Ad serving rates

We have looked at each of the stakeholder groups associated with an online advertising campaign and some of the technologies and processes they use in the lead up to the campaign setup, most of which occurs outside of the third party ad server. The actual process of setting up the advertising campaign and distributing the ads and creative through to the Publishers is done in a channel by channel process by logging directly into the ad server interface. But what does it mean to setup a campaign channel by channel?

Channels evolve, change and mature over time. There has been an effort on the side of the technology providers and Agencies to try and identify the properties which distinguish channels from one another and although on the surface they seem to be different, the underlying tracking technology is essentially the same. Essentially there are three types of basic tracking technology which allow for reporting in the ad server, these are impression tracking, click tracking and interaction tracking. In various combinations and implementations these trackers can be used to track advertising in all media channels but it is the properties of these channels which ad to the complexity behind the setup. The rule is that the more complex the Media Agency would like to have their data reported by the ad server, the more complex the setup needs to be.

With these complex insights it becomes easier to overcome new tracking challenges to feed advertising campaigns with more data to attain higher ROI. The features available to accomplish precision tracking in the ad server are now highly evolved and allow for vast changes from a single switch or adaptations to tracking in channels that might spring up literally overnight.

Which channels does an Advertiser need to run a digital ad campaign across to be successful? The quick answer is that there is no magic mix. Channels which have yielded the best results in the past become familiar territory to planners and buyers and are over-used, driving up the price of media (such as in the Search Channel). Constant channel experimentation is recommended since the right mix will requires an adaptable strategy.

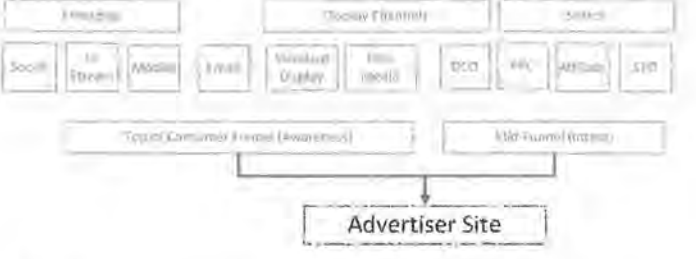


Figure 2.8.6 *: A reminder of the current state of play for the channels available to media planners and buyers on which to plan and buy inventory and drive the target audience down the purchase funnel. This example diagram is a guide rather than a rule.*

How does the end to end Campaign Management workflow work?

To review: the overall campaign workflow looks like this:

1. the Agency or Advertiser will typically make a deal with one or more Publishers for their inventory
2. the Creative Agency will build the ads.

3. The Agency/Advertiser will then take the reins themselves to manage the campaign, channel by channel.

It is on this last item that the third party ad server comes into its own; carefully tracking the whole ad campaign and with it tracking potential customers as well as delivering the ads to user machines, churning out the tracking data into reports (to understand performance) and then giving the Agency the option to adjust or optimize the campaign setup for the best return. This process is cyclical throughout the lifetime of the campaign and then the same lessons can be applied for future campaigns until the Agency/Advertiser treats the process as a KPI hitting production line, and great expertise can be born out of rinse and repeat.

Apparent exponential technology growth means that this production line changes dramatically every three years or so and therefore a hunger for learning and adaption (trial and error) makes for the most successful Advertising campaigns and the best digital media Agencies.

Billing, Electronic IO and Ad serving rates

Third party ad serving is a cost that sits on top of costs for the media space and although it is essential to keep costs down for every campaign, selecting a new ad server is not something that typically happens on campaign by campaign basis. Getting the best deal on "ad serving rates" is understandably important for all parties that have to bear the cost, but ad serving rates themselves are contractual and reviewed over a period of a few years. Ad serving rates are usually set a price on the delivery ads across the ad server reporting metric "impressions". However, compared to media costs, ad serving costs are significantly lower and for the most basic ad serving and tracking should very rarely take a sizable percentage of total campaign spend. Bolting on other technologies to improve insights, performance and to lower costs (such as DCO) will raise costs higher, but then the gains are

significant. Media Agencies managing campaigns on behalf of Advertisers sometimes pass the cost of ad serving directly along to the Advertiser but this is not a rule for all Media Agencies so it pays to shop around and ask questions about ad serving rates when an Advertiser is selecting a Media Agency to work with long term.

Now that the bills are out of the way, it is time to introduce the concept of "trafficking".

Chapter 3

Ad serving Operations

Chapter 3
Ad serving Operations

- Trafficking
- The third party ad server user Interface
- User Access for campaign setup/ Permissions
    - Using multiple third party ad-servers (double-tracking)
    - *Figure 3.0.1*
    - *Figure 3.0.2*
    - The golden rule of ad serving
- How does the trafficker check the ads are working? (Setup verification using an HTTP tracer)

Trafficking

The campaign setup process is referred to as "Trafficking" and is undertaken by either:

- The Media Agency
- The Advertiser
- The third party ad server support team (outsourced)
- A specialist trafficking company (outsourced)

The emergence of specialist trafficking companies is not a new concept in the world of ad serving. Big players include *Theorem*, *Traffic Rich*, and *Operative*. At these organisations large numbers of technically trained staff (usually based in countries where staff can be hired at a low cost) will complete the trafficking process according to the media plan supplied by the Media Agency. These companies provide support for most of the known ad servers and can be essential for large Advertiser marketing teams who do not employ an Agency and do not have the resources to complete trafficking themselves.

The speed of the trafficking process depends on the complexity of the campaign and the experience of the "Trafficker" (the person doing the trafficking). Media plans supplied to traffickers on a Friday afternoon for a campaign live date of the weekend all too commonly experience a delay owing in part to problems such as incomplete Creative Agencies supplied creative or problems meeting Publishers specifications.

The third party ad server user Interface

The rest of this Chapter will make reference to the Trafficker using the third party ad server to setup and control the tracking of the digital ad campaign. The third party ad server has a universal user interface which can be accessed on any computer connected to the internet. Once logged into this interface the trafficker can perform the role necessary to conduct these tasks. Although the interface itself differs in look, feel and design between the third party ad servers on the market, the features available in the tracking setup for each channel are almost identical. Where such features are unique to a particular ad server, this book will make reference to such facts and indicate which third party ad server the feature is unique to. As iterated in the introduction to this book, some of these features may be universally available in all ad servers or adapted by the time this book goes to print or in the future. It is important to check with third party ad server technology companies to get the most up to date information about feature availability.

User Access for campaign setup/ Permissions

A trafficker has the appropriate access levels inside the third party ad server interface to make changes to all aspects of the campaign. It is advised for Agencies and Advertisers to be fully aware of the trafficker's access permissions before they embark on the trafficking process. It is important to be aware of what parts of the setup could be subject to human error, if tasks are completed inadequately. It is for this reason that these technologies have very specific levels of enabler permissions in

place, so that some settings which may be tied directly to a fee or charge, are out of reach of inexperienced traffickers.

Now before we explore into the intricacy of each setup by channel there are two topics to touch on that apply to trafficking regardless of which channel is being tracked:

1. Double Tracking
2. Setup Verification of live ads

1. **Using multiple third party ad-servers (double-tracking)**
   Sometimes Advertisers will employ multiple Agencies to work on a campaign by dividing the work up by the channel that they belong to (for instance some Agencies might be specialists in running Search campaigns where others usually manage Display). When these Agencies are using different third party ad servers or site serving through Publishers directly, it is important to ensure that each ad server (or the main ad server used for reporting) can track all of the online activity as clearly as possible to prevent the duplication of audiences and users.

   What this means in practice is that a third party ad server will drop a third party cookie on a user when it serves them an ad (see Chapter 1). If two separate Media Agencies are reporting their campaign results back to the same Advertiser from two different ad servers, it will not be clear if they are describing the behaviour of the same individual. Furthermore it will be very difficult to describe how a user came into contact with both sets of activity over the same time period in order to distinguish which ad or Media Agency was responsible for pushing the user further down the purchase funnel.

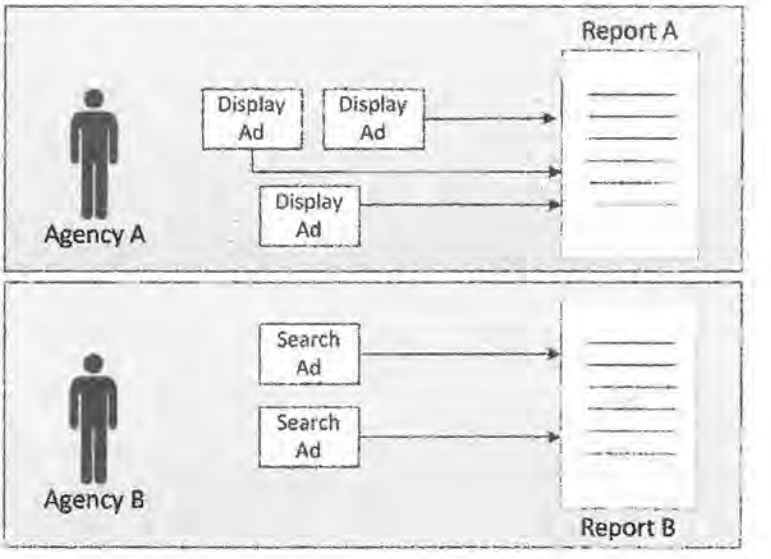


***Figure 3.0.1**: The problem at hand; working with two separate Media Agencies or at least two separate ad serving systems means that ads existing outside of the system are essentially invisible to that Agency and a true picture in single report, with a consistent methodology, does not exist. Building the results from multiple system reports is extremely challenging.*

Essentially the activity tracked in ad server A is invisible to the activity tracked in ad server B and visa-versa. On several occasions I have sat down with puzzled media planners who cannot understand why the search channel activity is being controlled by their sister media Agency, is not showing up in their standard third party ad server campaign reports. The answer is of course that the sister Agency is using a different third party ad server (or a section of the same third party ad server which is cut off from the main ad campaign tracking).

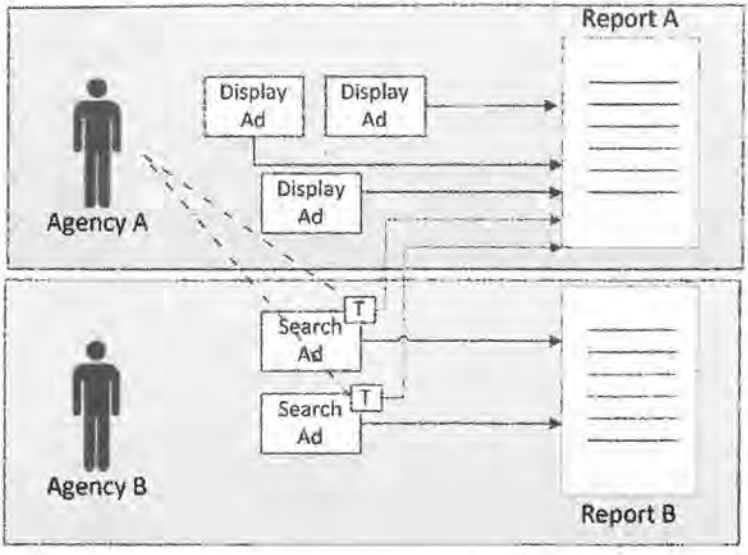


***Figure 3.0.2**: The solution: double-tracking. Cheap tracking pixel technology is generated and subsequently implemented to reveal the invisible ad activity to one of the Agencies(Agency A) who will now be in a position to represent all activity in a single report (Report A).*

To resolve this problem the traffickers will need to make sure that the ad server from which the final reports are being generated, creates extra tracking code to go into the activity being completed by the secondary ad server.

This way all the activity is tracked together (because embedded tracking code from the primary ad server in the second ad server reveals those ads so that they are recognised by the primary ad server as native ads belonging to the ad campaign).

The same should be done for Publisher served creative (the code needs to be created in the third party ad server and added to the ads when they deliver form the Publisher ad server if site-served). To allow for this tracking, the trafficker

creates a series of tracking pixels in the ad server (a process we will cover in more depth in the next section). This whole process makes de-duplicated reporting available (also called "De-duplication") but there is an ad serving cost to tracking things in two systems referred to as "double-tracking".

De-duplication is a golden rule for successful ad serving; always making sure that the entire audience for all ads belonging to an Advertiser, are tracked in a single system.

Double-tracking is a more advanced concept but an everyday reality for Advertisers who have a lot of digital activity to run and a need to bring their reporting data together to make sense of it. This is one of the advantages of using a single system for data management, although a single third party ad server does the trick, new technologies are now emerging that appear to facilitate this need such as Data Management Platforms (DMPs). This issue with such a method is that the tracking methodology between two wholly different ad serving systems will cause unseen errors in the collection and measurement of the data. This is much more advanced so we will save further discussion about DMPs until Chapter 8 - Programmatic.

2. **How does the trafficker check the ads are working? (Setup verification using an HTTP tracer)**
It is important to know once the trafficking process is complete that the ads are in place and being served as expected on Publisher pages. This can be done without having to access the reporting interface in the third party ad server and for peace of mind some Advertiser will want to be able to perform this check themselves.

An HTTP Tracer is a small program that runs alongside the browser and monitors the packets of data travelling to and from the user's browser. This is enormously useful when checking to see if code has been implemented on a page correctly. By loading the publisher page up where the ad is displayed and having the tracer open the server calls can be viewed in real time, demonstrating that tracking is indeed in place.

Most tracers are free. *Fiddler* (a standalone program) and *Firebug* (a *Firefox* add-on) are popular choices. As well as seeing if ads are clicking, redirecting and loading up parameters or other files, the tracers can note if event handlers are being fired for Rich Media ads and display variables that are being passed on to the server (for further details about setup verification please see Chapter 3 - Display ad serving).

On day one of campaign launch, the trafficker can utilise monitoring tools and then pull reports from the third party ad server to check on delivery, performance and total sales or conversions (see Chapter 4).

These are two generalised (non-channel specific) considerations for the seasoned trafficker, but the full display trafficking process lies ahead.

Chapter 3 - Section 2
Display ad serving operations

- Day in the life of the trafficker
- **Foundation Level trafficking for standard display**
    - Creating a place in the ad server to setup the campaign
        - Third party ad server hierarchy
        - *Figure 3.1*
        - Settings at Agency and Advertiser levels - Silo Capabilities
        - *Figure 3.2*
        - Settings at Agency and Advertiser levels - Sharing Capabilities
    - Structuring the campaign setup to replicate the Publishers on the media plan
        - Media Plan
        - *Figure 3.2.1*
        - Trafficking Sheets for outsource trafficking
        - Sites
    - Identify the formats of the ads and creating the placements
        - Ad Formats
            - Banner
            - Leaderboard
            - Skyscraper
            - MPU
            - Roadblock and Skins
        - Placements and Packages
        - Cost Packages
    - Assign the creative and setting rotation controls
    - Naming Conventions
        - Ad & Creative
    - Automated QA
    - Ad Preview
    - Adding the click-through URL: clickTag

        - 302 Redirects
        - Adding multiple click-through URLs
        - Ad classifications / Custom Categories and searching
- Ad Duplication
    - *Figure 3.2.2*
    - *Figure 3.2.3*
- Readying the Ads for Assignment
    - Backup gif
- Possible Placement Assignments
    - *Figure 3.2.4*
- Ad assignment
    - Multiple ads to a single placement
    - A single ad to multiple placements
    - A single ad to a single placement
    - *Figure 3.2.5*
- Adjusting the Ad Controls
    - Frequency Capping
    - Ad Rotation
    - Sequencing
    - Weighting
    - Ad Prioritization
    - Time Based
- Post Cap Gif
- The two levels of control
    - *Figure 3.2.6*
- Trafficking the Tags to the Publishers
    - Ad Tags and Tag Formats
        - *Figure 3.2.7*
    - Flighting
    - Thoroughly Checking an Ad is live
- **Intermediate level trafficking for standard display**
- How does a trafficker manage campaigns with hundreds of Ads?
    - Bulk trafficking through Excel templates
    - Bulk trafficking with Excel as the UI
        - Automated trafficking setup via ad server API

        - *Figure 3.2.7.1*
    - The JavaScript carriage capability
    - Survey vendor Code and third party scripts-in ads
        - *Figure 3.2.8*
    - Tracking but not serving
        - Is it possible that a Publisher cannot support third party ad serving?
    - *Figure 3.2.9*
    - A note on Click Commands
- Ad Replication
    - Ad Delivery and Replication
        - Server-Side Caching
    - What exactly is a tag?
- Making a bigger impact with Display Campaigns

This section runs through the trafficking process from end to end, using the Display Channel as an example. Campaigns for Mobile, Instream, Rich Media and DCO generally follow the same workflow (i.e. the creation of the Media Plan, the adding of Sites, the creation of Placements, the assignment of Ads, the additional of Ad Controls or Weighting rules and the sending of the ad tags to the Publisher). As we are focussing on the Display Channel, and the majority of the legwork is done by the trafficker, it is this role that will undertake the majority of tasks that follow.

Day in the life of the trafficker

Prior to the trafficking process, the media plan makes its way around the Media Agency, before being handed over to the designated trafficker. In an ideal world when this document is completed: the Publishers sites are listed, the number of booked impressions is highlighted against each one and the creative names are filled in to an agreed convention. Media Agencies are rarely able to bring together so much concise information in time, and more often than not will pass along the media plan to the

trafficker incomplete. Media plans should come as close to a standardised document as possible (being the same document layout used in all Agencies around the world) but currently media plans come in lots of shapes and sizes. For this reason trafficking companies and some ad servers issue a template called a "trafficking sheet". The trafficking sheet is just a media plan, filled out in a specific layout and the benefit is that this allows for high speed human or machine processing.

The trafficker will work from the layout in trafficking sheet, attempting to replicate the required setup in the user interface of the third party ad server. Information missing from such a document slows down the trafficking process and ultimately the campaign. Adhering to the strict instructions coupled with the trafficking sheet to aid its completion, should be followed to the letter to keep the correct pace for the digital ad campaign delivery timeline.

The trafficker will keep to the following workflow:

Foundation-level Trafficking for Standard Display

1. Creating a place in the ad server to setup the campaign
2. Structuring the campaign setup to replicate the Publishers on the media plan
3. Identify the formats of the ads from the supplied plan and create the placements to match
4. Assign the creative to these placements and set rotation controls
5. Traffic the tags

The rest of this section will focus on this simple workflow and then look at the trafficking of more advanced media plans, showing how the trafficker can take advantage of bulk trafficking technology, working directly from excel.

Intermediate-level Trafficking for Standard Display

6. Bulk trafficking and 4th party tracking
7. Ad Replication
8. Making a bigger impact with Display Campaigns

1. Creating a place in the ad server to setup the campaign

The success of a campaign and the data that the third party ad server collects, not only depends on a clean and effective trafficking process and campaign setup, but also on considerations made by the trafficking team long before a campaign begins. These considerations include the organisation of all campaigns at a higher level in what is termed "the ad server hierarchy". Planning and implementing an optimum hierarchy setup will save hours of work later on.

The third party ad server hierarchy

The hierarchy of the third party ad server is regrettably not uniform across all third party ad servers and is a reason that many Media Agencies have a preferred third party ad serving provider. Switching to another ad server can mean a re-education into how the ad server hierarchy is constructed.

To review from the beginning of this book; when a trafficker first logs into the user interface of the third party ad server they will note that the whole interface is divided into Campaign Setup versus Reporting/Analytics. Some ad servers issue different access credentials to the trafficker to divide these areas up, restricting access to those that do not

need to see particular parts of the interface. Should the trafficker wish to begin the process of Campaign setup, they navigate to the campaign setup section.

The first thing to note about campaign setup is that access is limited to the Advertisers and Campaigns that the trafficker has been granted permissions to both read and append. To coin a phrase from the world of computer programming, the trafficker is said to have "read and write permissions". With the correct permissions in hand, the hierarchy that the trafficker initially arrives at has many levels. Each of these levels serves a purpose.

DoubleClick (DFA) hierarchy is structured so that the highest level is called "Network", the next level down is called "Sub Network". The traditional *Atlas* hierarchy begins with the level "Company" followed by a level called "Branch" before moving down to "Client" and then to "Advertiser".

Mediamind offers an "Advertiser" and "Advertiser-Brand" level at the top of the hierarchy. Each of these folder-like entities has a rather unique set of rules. The more folders there are in a hierarchy (both horizontally - across to multiple Advertisers and vertically, down to campaign level), the more situations that a third party ad server can cater for, these include very advanced use cases.

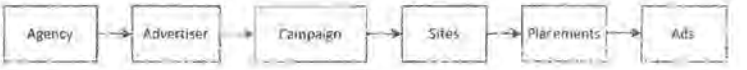


Figure 3.1 *The trafficker initially logs into the third party ad serving campaign management and setup interface and drills down into the standardised hierarchy. The very top levels of the hierarchy are commonly hidden from the trafficker so the trafficker appears to begin the journey in a folder for their own Agency displaying all the Advertisers that they manage. In turn the Advertiser folder contains all of the Campaigns that have been setup or are being managed under that Advertiser. Please note that we are focussing on the setup for Standard Display campaigns in this section. Hierarchy navigation is not completely consistent for all channels, most notably the arrangement is different for the "Search" channel..*

Settings at Agency and Advertiser levels - Silo Capabilities

The top level entities in the third party ad server interface, offer their own settings which can be applied to the folders and entities that they contain. This is useful for applying bulk settings to the hierarchy that sits below. Examples of the unique capabilities of these top level entities are as follows:

a. Allowing for the classification of the whole Advertiser into a vertical definition such as "Automobile" or "banking" (this information is fed into benchmark reporting and publicly published by the ad server on a rolling cadence - See Chapter 5 - Reporting)

b. Billing Settings (so that all activity associated with the account can be captured by the ad server billing systems and the ad server costs can be invoiced to the paying client)

c. Contact Settings (usually Publisher contact details for the distribution of ad tags - see later in this section under 'Flighting')

d. Bulk application of third party tracking or third party applications (such as the deployment of the Ad Choices badge to all ads - more detail of which is available in Chapter 7 - Retargeting)

e. Advertiser Universal settings such as Verification settings, Time Zone or location settings

This last item applies a rule to the Advertiser as a whole to establish how data should be collected, stored and processed by the ad server. The presence of these settings constricts the data collection process which informs the reporting system. In essence they act as filters and have a big impact on reporting when adjusted. Given the more complex nature of these settings they are discussed in more detail in Chapter 5 - reporting, and are omitted from this section.

Ad server hierarchies that silo their data and tracking are more common place than those that offer the alternative sharing model.

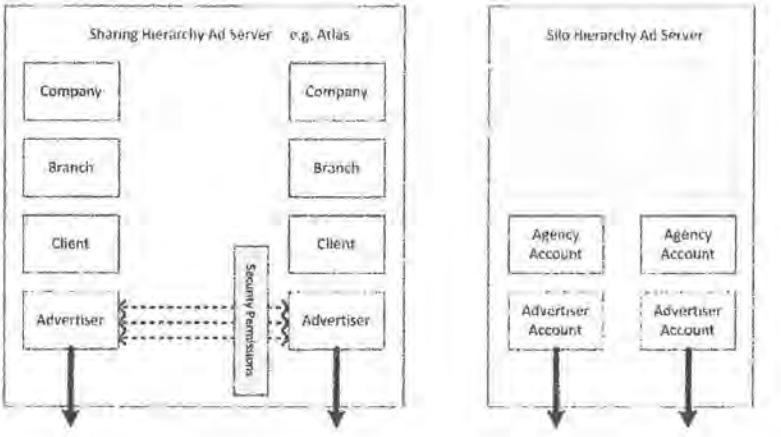


Figure 3.2 *The two types of account structure or hierarchy are shown: The Silo model and the Sharing Model. Siloed models; although more simple, are a little restrictive when big Advertisers need to extend the influence, control and counting methodology across to other Advertisers and Agencies that may be working together under a joint venture.*

Settings at Agency and Advertiser levels - Sharing Capabilities
Under the sharing model, the top of the hierarchy can be customised to share information with other top level entities. This is a feature common place in the original *Atlas* third party ad server structure and allows for campaign-collaboration between multiple Agencies and Advertisers. The most common settings which are shared utilising this 'bridge' are:

- Conversion Tags (see Chapter 4 for description; this is an alternative method of data-sharing to the more popular "piggybacking" method)

- Double Counting (Duplication and De-duplication settings) The presence of this setting allows for the tracking of channels like Natural Search (see Chapter 3 - Search)

- Threshold and Time-zone Settings (porting Silo settings across to other Advertisers that may benefit from standardising a particular setting amongst them).

Control of these shared settings and entities is given to the source entity/account. It is seen as a benefit to be able to use a single account to control settings in a remote and separate hierarchy. Issues can occur where a trafficker amends a particular setting, unaware that the adjustment affects other (seemingly disconnected) accounts. It is an important feature of the Ad Server user interface therefore, to clearly broadcast to the trafficker if any part of the hierarchy is being shared to other accounts, and precisely where that is occurring.

2. Structuring the campaign setup to replicate the Publishers on the media plan

After the trafficker has logged into the ad server user interface (and is presented with a list of all the Advertisers that they have been granted access to), the correct Advertiser folder is selected and clicked on, bringing the trafficker to a list of Campaigns that are stored under that Advertiser.

In these early rafters of the ad server hierarchy, the trafficker will not have the required permissions to created new Advertisers or Agency's. This work must be undertaken by the ad server support team (to prevent hierarchies from being crowded with new Advertisers). Owing to this, Advertisers new to Agencies or new to third party ad servers take extra days to be setup for the first time. Only once an Advertiser entity is in place, can campaigns be created within it.

Media Plan

As covered in Chapter 2, the media plan is a manufactured documentation keeping the record of all Publisher bookings

for multiple Publishers within a single campaign. From the Media Agency's perspective, the Media Plan is that documentation; however "media plan" is also a third party ad server hierarchy entity in its own right. This means that the Media Agency version of the media plan needs to be processed by the trafficker (or adapted into a trafficking sheet) so that the plan is fully replicated in the ad server.

The trafficking sheet, is predominantly in a spread sheet format with each row representing the relationship between a Publisher website ad location/spot (called a "placement") and the creative/ad that is to appear in that spot. The media plan also looks much like a project plan or Gantt chart as it features a timeline which clearly demonstrates how long the ad will be assigned to the placement and displayed on the Publisher site for.



Figure 3.2.1 *A example of a media plan.*

Since physical media plans rarely look like the end result in the ad server user interface, some Media Agencies are asked to convert their media plan documentation into the "trafficking sheet" layout or template referenced at the beginning of this section. Such sheets are particularly useful because they can be fed into automated systems, which can generate the required setup without utilising the time and resource of a trafficker. Some in-agency planning tools covered off in Chapter 2 are capable of undertaking this role, side stepping the need for as much trafficking headcount on the Media Agency side. These systems (such as *Adazzle*) have

an API connection with the third party ad server in order to enable such automated processes. (So as not to diverge from the core trafficking process this is explained at the end of this section under this section in *Automated and Bulk trafficking*.)

Trafficking Sheets for outsource trafficking

When outsourcing trafficking to an external trafficking company, supplying a completed trafficking sheet will be more assistive than supplying the original media plan. This is because the traffickers at these organisations will be processing the media plans of a large number of Agencies and Advertisers in a day. This requires uniformity to complete the work effectively. Agencies that do not employ traffickers and rely on account teams to construct the media plan, are responsible for completing the trafficking sheet in full or face the rejection of the documentation by the outsourced trafficker. Incomplete media plans and trafficking sheets continue to play havoc with digital campaign delivery times where such rules are not strictly enforced and rejection, due to incomplete sheets, is commonplace.

Sites

Since each row of the trafficking sheet references a particular Publisher website, this information must be replicated in the campaign setup. The trafficker begins by creating a new media plan entity and then drilling into that entity to add a new site to the media plan. A master list of all known Publisher sites appears on the screen and the trafficker makes the correct selection.

As rosy as this may sound, the master-list of all known Publisher sites is a centralised database of the third party ad server that requires regular maintenance. A universal agreement between all traffickers, ad server users and Agencies exists through the enforcement of anti-replication

rules in the ad server so that only one of each site appears on the list (those sites that appear to be absent can be added upon request to the ad server support team, but this process can take a day or more and should be factored into the timelines for the campaign setup).

The duplication which realistically appears in these databases (because each master list of sites database is unique to each third party ad server) is chaotic. Some popular sites have so many names to represent them it is impossible for the trafficker to decide which one to choose. The best practice? To pick a single name and continue to use that same name every time media is booked on and ads trafficked to that website.

Why can't the site master-list be unique to each account, Agency or Advertiser? If it was, there would be far less clutter, the data would be far "cleaner" and it would take less time to setup and traffic a campaign. The reason to retain a single consolidated list is because Publishers also access the third party ad server. When Publishers access the third party ad server they see a list of all of their sites and the activity being run across them by all Agencies and Advertisers. Traffickers that select the wrong site when setting up the media plan (despite the fact that the name might look right) are inadvertently messing up Publisher-reporting from the ad server interface (as seen from a different user perspective). If the trafficker is unclear about exactly which name to use when setting up the Media plan they should contact the Publisher for that site and check in with them to see which site names listed in the ad server interface match with the media bought and listed on the plan.

A typical media plan contains upwards of one Publisher website, some big campaigns feature twenty or more sites

and may also include Publisher networks (which also are selectable as ad server "site" entities when setting up the media plan). Once the Publisher and sites are selected, the foundations of the media plan are complete. The next step is to copy the trafficking sheet further which will show that each site has at least one corresponding "placement".

Placements are dependent on one piece of information above all others and that is the dimensions. Without the correct dimensions declared, ads will always be displayed incorrectly; the placement must be considered to be the frame of the ad and the ad must always fit into the frame. To get the dimensions of the placement right, the site needs to be matched up with what the Creative Agency developed. The collective name for all the creative produced to standardised sizes is called "formats" and the correct formats need identifying to proceed with an accurate trafficking setup.

3. Identify the formats of the ads and creating the placements

When Publishers decide to sell ad space they take a long hard look at the their website and establish the opportunity to package the available inventory to sell to Advertisers of Agencies. By limiting the dimensions of the ads so that across a site there are thousands of ad spots available with dimensions of type one and thousands available by dimensions of type two, there are only two total sets of dimensions and therefore only two pieces of creative that the Creative Agency needs to build. These two sizes, are referenced by the Publisher as its "available formats". Dimensions however are just a single property of what defines a format. The most basic format is an ad with dimensions 300 pixels wide by 200 pixels high but more complex properties to these dimensions (such as the ability

for the ad to cast the whole screen to one side) changes the name of the format from the industry recognised "MPU" to the industry recognised "sidekick". It is the standardisation of formats for all Publishers right across the web which has made ad campaigns manageable and scalable.

Ad formats

The web is full of ad space but in major markets like the UK, the *IAB* has defined most of these spaces as being capable of occupation by "standard formats" which are represented by various creative dimensions as well as other properties as discussed. This standardisation has helped Creative Agencies to build ads faster (and more effectively). It has also helped Publishers easily monetise their inventory offerings in the wider marketplace. Ads built to the IAB specifications and dimensions of a standardised format can easily be trafficked to all sites supporting IAB approved formats, with the standardisation escalating the overall process.

As previously mentioned in Chapter 2- Creative Agencies, the *IAB* also issues guidelines for ad creation and acceptance, to maintain a good user experience on Publisher sites:

- Standard flash ads for instance are advised not to "loop" their animation more than three times (repetitive messaging inside a single piece of creative) which has been shown to deter the positive engagement of a user with a brand.
- For ads that might contain video, it is advised to keep the audio of the video muted when the ad loads, even if the Publisher permits the video to begin playing through without it having been initiated by a user (sometimes referred to as an "auto-play" as opposed to a "user initiated video play")

- High contrast colour usage and flashing imagery which may pose a risk to users with photo sensitive epilepsy have vanished
- And perhaps most famously is the strict guideline against the use of JavaScript "pop-ups" and "pop-under's" which plagued the early web and had a significantly negative impact on user experience.

Although the IAB do not get involved in the ad development process with the Creative Agencies; the Publishers themselves have adopted these guidelines into their own strict specifications so that ads containing too many loops for instance are rejected before the campaign can go live.

The use of the term "format" can be understood by considering a "pop-up" ad format. The properties of the ad are almost the same, such that when the "pop-up" format is run on any site, everyone knows what it is and how to get rid of it. This is the ad format standardisation that the IAB have been largely responsible for; bringing order to an immature internet circa 2004. Format names are one of the most important concepts to grasp before the trafficking process can begin because they are the content of the media plan or "trafficking sheet".

Banner

The banner format or 468 x 60 was very popular when it was originally coined as a standard format but today, as screens have grown in size, this format has become less and less popular because the ad itself is really too small to push a creative message and remain impactful, on a computer and remains too large for more portable devices.

Leaderboard
The leaderboard is a 728 pixels by 90 pixels ad format, and has taken the limelight away from the banner format as it increased in popularity, simply because there is more space to work with. Some Publishers offer a super leaderboard format which will typically be more than 900 pixels across and is sometimes referred to as the "masthead".

Skyscraper
The skyscraper (120 pixels by 600 pixels) remains a very popular ad size on Publisher sites although it is beginning to be overtaken in popularity by the larger "super skyscraper" at 160x600. Once again because screen sizes in non-mobile and non-tablet devices are growing, so the super skyscraper just presents more information.

MPU
The MPU (mid page unit) is the squarer sized ad format at 300 pixels by 250 pixels. Still a very popular format, the MPU is very commonly used as a space to host Rich media video as the square size lends itself nicely to building a video player. MPUs are commonly more expensive spaces as well, as they are more likely to appear closer to the top of page content as a single unit, than the Leaderboard and Skyscraper formats (which often appear multiple times on a page and are more likely to drop below the fold of the browser window).

Roadblock and Skins
A Roadblock is the name given when a Publisher page has all its ad spots taken over by the same Advertiser and ad campaign. Roadblocks consisting of two or more ads can be synced together using Rich Media (see Sync Ads), but even when unsynchronised, still offer more impact by repeating the same message multiple times on the same

page. Sometimes web pages also centre all of the content of the page (more common with the major Publishers) which leave the sides or "gutters" of the page available for advertising. A "skin" is giant image or collection of images that cover these areas: also referred to as the "Wallpaper."

Perhaps we have taken something of a detour from the world of trafficking into the universe of formats but to bring the story home, each row of the trafficking sheet or media plan will list out a *placement* of a particular size, (this is the housing for the format of the creative it will be paired with and display, when the ad loads). The trafficker selects the site in the media plan in the ad server UI and is taken to a new blank screen which the trafficker will populate with all the placements under that site which appear on the media plan.

Placement & Package

Inside each site or Publisher entity in the media plan the trafficker creates a series of "placements" which will match the corresponding ad formats purchased from the Publishers in the plan. The trafficker can check the physical Publisher campaign IO to check the setup correctly matches with what was purchased and what appears on the Media plan.

Placements belonging to Publishers at the media plan level in the third party ad server interface, can be grouped together by the trafficker into "Packages" if required, so that some bulk settings can easily be applied to a large number of placements. These settings usually refer to the monetary value of the media cost of each of the placements which are *labelled* as "Costs" in the ad server. To allow for added flexibility, packages support different cost models and are the reason most traffickers decide to keep a pocket

calculator handy on their desks during this portion of the trafficking process. For this reason Packages are sometimes referred to as "cost packages".

Cost Packages

The inputted information in the cost fields in a cost package is for reference purposes only and does not carry through to the Publishers themselves. This means that if the trafficker were to set the costs of a placement and then amend them later - the agreed IO cost and contract with the Publisher is not changed or broken: this information is never conveyed to the Publisher themselves. Instead it is used to calculate extra cost fields in the ad server reporting, based on information like delivery and performance statistics. For instance, the cost field can be aggregated across all of an Advertiser's campaigns running over a given period and then broken down by country to calculate the cost amount being spent in all countries across all campaigns at a set period in time. In this example the Agency or Advertiser could pull this information from the reporting section of the ad server interface (more information in Chapter 5 - Reporting)

As cost information additions are extra work for the trafficker they often omit this but then find the need to - rebuild reports manually later on to calculate these post-campaign costs. It is highly recommended therefore to include these figures during the trafficking process so that the ad server does the leg work of calculation. Some ad servers also have the option of displaying ad serving costs and Agency mark-up costs to reflect the true cost of each package to the Advertiser.

For some ad servers, placements cannot be created without first creating a package to put them in, and the media plan screen looks like a compressed version of the trafficking sheet. Here, placements are hidden inside the packages as

folders and it is the packages which are assigned to the sites. Once package settings are created they are rarely changed, so packages are often hidden to the trafficker in the trafficking screens after creation (to remove clutter on the screen).

Once Publishers, Packages and Placements have been setup in the ad server interface by the trafficker, the next step is to match up the creative with the ads.

4. Assign the creative and setting rotation controls

In parallel to the trafficker setting up the campaign in the ad serving interface, the Creative Agency connects to the ad server via the supplied components in the flash authoring environment. It is here the Creative Agency can name the files that the developer supplies. A successful upload will return a confirmation message inside Flash with a unique identifier for each uploaded file.

These identifiers can then be sent on to the trafficker, who is able to retrieve the creative. In *DG/Mediamind* for instance these identifiers are called "Ad ID's". Ad ID's themselves are a poor indicator for what a creative actually looked like or where it came from, so the file-name supplied by the developer is quite significant. Despite the fact that the creative developer will supply the creative with a certain name, this name can be altered by the trafficker at the start of the trafficking process, so as to filter down effectively to reporting later on.

Naming Conventions

Creative naming conventions (as in the filenames given to the Flash creative and GIFs) can mean the difference between efficient and inefficient trafficking. Creative naming

conventions should be agreed between the Creative Agency and Media Agency before the ads are created, so that when they are uploaded, the context of the files supplied is perfectly clear.

Some information that the name could be made up of might include;

- the creation date
- a brief description of the creative
- the campaign name
- the Advertiser name
- the name of the Creative Agency developer
- information about the intended location of the final ad (if known) such as placement or site information
- an indicator of some of the functionality of the creative such as "expand" for a Rich media expandable ad
- an indicator of the ad's intend KPI (such as "service signup" or "brand awareness")

Although most of these data points reference things that are added later in the campaign, such information really helps the trafficker working in the third party ad server to match the ad to the correct site, ad spot and creative folder.

Down the road the creative naming convention is also a lifesaver in reporting where the "creative name" field can be added to any kind of report revealing information about an ad that had been misplaced, or even a creative theme that was more successful than others. In many instances, the creative name can be used as quick fix where perhaps a third party ad server has reporting limitations (the naming convention is used to 'carry' creative information through to data analysts to improve campaign optimization).

To demonstrate the effectiveness of this, a creative name can be transformed from:

| ***Creative3_1108_dave***
| to
| ***Unilever_Lynx_Axe_brandawareness_summer2015_akqa_themealpha***

Any analyst reading a report containing any information about this creative will now at least have an indicator about the Advertiser and brand it is referring to owing to the references in the naming convention.

Ad & Creative

To use Creative in a campaign it must turned into the entity referred to as an "Ad". Some ad servers do this automatically, others require a couple of extra steps on the part of the trafficker to create the ads. This will be the first time that the trafficker has exposure to how the ad will look and behave, and as such, the trafficker can conduct some basic testing to ensure that the Creative Agency has built the ads to the requirements of the brief.

Automated QA

Some ad servers automatically process the creative when it arrives in from the Creative Agency. These checks for Quality Assurance (QA) check a short list of creative properties that should be present in all creative supplied, so these are standardised properties and not checks for unique functionality. Checks will include locating the presence of the clickTag, checking the creative for malicious code and ensuring that the click portions of the ad actually click when such user action is taken. Once ads have passed this automated QA they can be used for further testing by the trafficker. As a side note: some ad servers are not using machines for this process but hire teams of individuals to manually check each and every creative that passes into the ad server interface, as such the initial QA can take up to

around twenty four hours before the rest of the trafficking process can be undertaken.

Ad Preview

The newly created ad can be viewed in the ad server interface as it would be displayed when live. This preview area comes close to simulating how the ad server will display that ad once it is live on the Publisher website. The preview screen allows the trafficker to visualise the finished ad for the first time, check that interactions are firing (see Chapter 3 - Rich media) and that the ad is behaving as expected. The ad can then be viewed against the backdrop of a live Publisher page. Finally, a copy of the preview can be sent to any other Advertiser or Agency contacts for final sign off (it is not common for the sign off to be a mandatory process as this has proven to slow down the trafficking process). This shared preview produces a URL which often remains live long after the campaign has ended so that it can be used as a reference to what the ad looked like. This same URL can be supplied to the Creative Agency who would store it for their own on-site creative archive records so that the creative could be re-viewed as an ad at any point in the future.

Once the trafficker is satisfied that the ad is behaving as expected it is time to add a location for the click.

The main difference between 'Creative' and 'Ad' is that a usable ad has a corresponding URL so that when the user clicks on the ad when it is on a Publisher page, the user we are redirected through to the destination designated by the trafficker.

Before moving forwards it is worth restarting this claim: A completed or "finalised" ad (as it is referred to in the *Atlas* third party ad server) is the marriage of the creative and the clickthrough URL. The trafficker of course will take the URL

from the one supplied on the media plan. This is why creative is not assigned directly to placements, since creative does not carry a click-through URL.

The recommendation for any Advertiser is that the landing page (that the clickthrough URL directs to), should in some way be contextually relevant to the user. Some traffickers work with Advertiser website management teams to ensure that the clickthrough URL is dynamically populated with a value that the user may have inputted into the ad (such as zip code). The corresponding Advertiser site is then setup to process this value and amend the landing page in some way, such as by using the zip code to conduct a search and then displaying the search results specific to that zip code as the landing page. These more relevant experiences are a campaign consideration which will undoubtedly show improvements in campaign performance if set up correctly.

Adding the click-through URL: clickTag

Remember the clickTag concept from Chapter 2 - Creative Agency? As a reminder, all creative needs to be embedded with the correct clickTag by the Creative Agency to be approved for use inside the third party ad server for the campaign.

The clickTag is actually a very simple piece of technology that prevents the Creative Agency from "hard-coding" the clickthrough URL into the creative itself. Now, when a user clicks on an ad, the clickTag in the flash file queries the ad server, sending the user to the complete URL inputted by the trafficker rather than any address inputted by the creative developer.

The purpose of removing the power of the click-through URL from the creative developer and placing it instead with the Media Agency trafficker means that, should the click

through URL change during the campaign, the SWF file produced by the Creative Agency does not need to be adapted and "re-published".

This time-saving functionality is one of the reasons a clickTag is so important and also means that creative can be developed and published even if the final click-through URL is not known until the day before campaign launch. It allows for the flexibility of changing a URL midway through campaign. The other primary reason for the clickTag is that without it, the 3d party ad server would not record a "click" when the user clicks on the ad - more on this in Chapter 5 - Reporting.

302 Redirects

The response delivered from the ad server back to the browser commonly carries the HTTP status code of 302 (most readers will be much more familiar with the code 404 from the same family indicating a missing web page). By comparison 302 is used to inform the user's browser that the web page has been temporarily moved. This instruction allows the ad server to use the 302 call to log an impression on the counting module component of the ad server while serving the creative or 1x1 from the content server or CDN (see Chapter 1 - Ad Serving).

Adding multiple click-through URLs

As a side-note, the big third party ad servers consider a standard flash ad to be one that contains a single piece of animation or flash file, with a single click-through URL to correspond with the click-able areas. In some instances there are:

1. multiple locations within an ad to click, directed to the same page (which is still considered a single click-through URL)

2. multiple locations within an ad to click, directed to different pages (multiple click-through URLs)

In the first instance all the clicks will be attributed to the same "clicks" bucket when reporting on a single ad, to break out which click-location inside the ad was generating the most clicks the second scenario has to be adopted (ii). For (ii) rich media technology is usually utilised so that the second click is actually a custom interaction and the supporting report either reports "other" clickthroughs as custom interactions or redefines them in a custom report as an additional set of clicks alongside the same ad.

Alongside the click-through URL, it is worth the trafficker adding extra contextual information to the ad so that the performance of that ad and others with the same properties can be generalised in the reporting interface later on.

Ad classifications / Custom Categories and searching

By converting a creative entity into an "ad" entity, new options become available, options which are available to all ads said to sit at the "ad-level" unique to the Advertiser account. One of the key options available at the ad level is the ability to add fields to the description of the 'ad' in the ad server. One of the most useful of these fields common to the big ad servers is called the "ad classification" field. This may not be a single field but multiple fields. Each ad classification field is designed to carry a new piece of information about the ad entity, with that ad, on its journey from creation through to reporting.

For instance "ad classification field 1" may contain a creative message type. If I am a big automobile Advertiser my possible creative message types might range between 1) Model awareness 2) Model interest 3) Model detail. Each

ad created from a creative can be classified into one of the three types. It is up to the trafficker during the trafficking process therefore to classify the ads by appending the text written in "ad classification field 1" to suit one of these options. When the time comes to report on the performance of the ads in the ad server analytics interface it will be clear by displaying this field, which creative message was most effective regardless of other factors such as the ad size, the Publisher or if the ad contained any video for example.

Current ad serving technology offers ad classification fields as open text fields that the trafficker is in charge of updating. It is important to offer this flexibility which has been used creatively by many of the industry's problem solvers to improve the quality of reporting.

Other fields which sit at the ad level include the length and width of the creative. In older ad servers these fields also need to be manually updated by the trafficker to reflect the correct size of the creative but some accurately detect the size and populate these fields automatically. As with the ad classification field, the length and width of the ad can be added to reports to determine if ads of a particular size are better performers.

There are a whole selection of other fields which can be populated manually at the ad level and it is this data inputting task for each ad that is responsible for taking up so much of the trafficker's time.

Now that the ad has been well described in the ad server interface, any member of the Media Agency team that might go into the interface will be shown information about that ad and find it far easier to identify. The next step will address how the ad will relate to the Publishers on the media plan. This rests on ensuring that there are enough ads to go round.

Ad Duplication

In order to keep campaign costs down a minimum amount of creative will be produced by the Creative Agency; after all, what is the point in paying the Agency for a creative for each Publisher on the media plan, if the ad to be displayed will look exactly the same on every site?

In such an instance the creative is first transformed into an ad entity by the trafficker and then replicated so that each ad assigned to each Publisher's placements is unique (even if the only visually unique thing about it is the ad ID, generated unique for each ad created). Unique ad assignment is important because if the unique ad ID is analysed by itself later on for its performance, the unique properties of its assignment will reveal where it sat in the media plan and therefore which Publisher the performance could be attributed to.

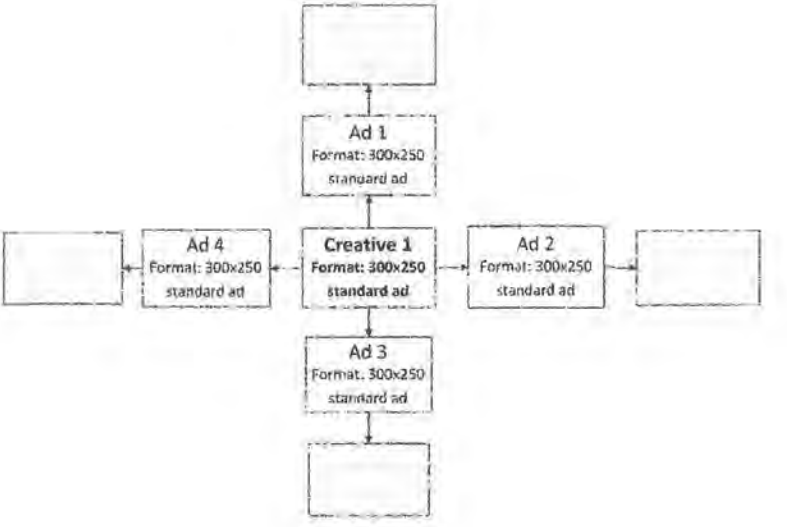


Figure 3.2.2 *making full use of an expensive Creative-Agency produced creative in a specific format. The Creative is simply duplicated as an Ad entity. In this image each ad is then unique to each Placement.*

In turn the trafficker is required to assign each ad to a single or multiple placements based on the information supplied by the Creative Agency (a creative name or ad ID) and the media plan in front of them.

It is not unheard of for the Creative Agency to side step any connection to the third party ad server entirely and just send the creative assets directly through the to the trafficker. The trafficker can log into the ad server interface and upload these files to an area within the Advertiser's "account" called the "creative library" at any time. This is the same location that the Creative Agency would upload their files to if they were using the automated upload facility available in the flash-based authoring component.

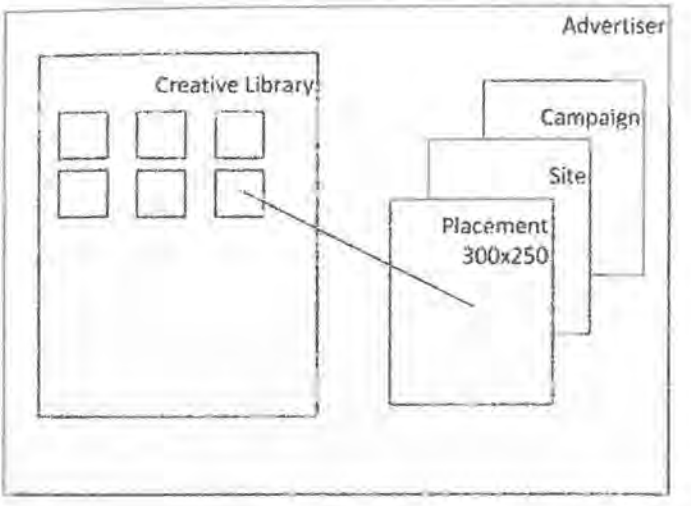


Figure 3.2.3 *The Creative library is present at the Advertiser level. Creative can be archived, arranged and stored in the creative library facility of the ad server. In the placement creation and ad assignment screens the trafficker choose the creative of the right dimensions and marry it with a corresponding placement.*

Readying the Ads for Assignment

Once ad entities have been created and the classification fields have been completed, the ads can be assigned to or coupled with a single or multiple placements. To enable the process of assignment, the placement needs to meet mandatory requirements:

- The ad has a corresponding click-through URL
- The ad is the same dimensions as the placement requires it to be (there are caveats to this for Rich Media ads)
- The ad has a fall-back display option, in case the user's browser cannot display a flash file.

The suggestion of a fall-back option is a new concept, but one worthy of further explanation:

Backup Gif

Standard ad serving and Rich Media offer a fail-safe, so that user machines or devices that do not have a flash player installed (or have an older version than the one required to play the creative) still see an ad. The fail-safe ad is a "backup gif" and must be supplied whenever a flash file is supplied by the Creative Agency. Since the third party ad server must always serve something if an ad request is received by the server, it will show the backup gif when the conditions of absent technology on the local user machine or browser are met (such as the absence of the flash player plugin).

Serving backup gifs is a more relied upon failsafe in recent times with the lack of flash support for mobile devices. In instances where a mobile-specific campaign is not setup, and a user browsers to a Publisher site intended for a PC user; an *iPad* for instance, fails to receive the flash file and a backup gif is displayed instead. Such failsafe or fall back technologies are to play a wider role in the development of ads for mobile and flash alternatives such as HTML5.

When the campaign goes live the placements should rarely be showing the backup gif since the whole point is to be showing the best possible type of creative to the end user. It is possible to monitor these occurrences from within the third party ad server reporting interface; occurrences might be confined to a single placement or Publisher which would require further investigation in conjunction with the Publisher ad ops team directly.

Possible Placement Assignments

As a rule, the assignment of ads to placements is not exclusive to the setup of 'an ad (with click through URL) with backup image'. There are other setups that can be assigned to Placements, the backup image can be assigned on its own without a better ad in place and the campaign can still go live.

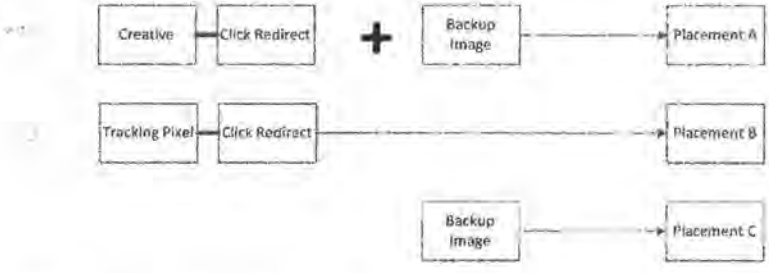


Figure 3.2.4 *So far we have discussed the setup shown below for Placement A (the ad to placement assignment for Ad 1). Placement B shows the setup for channels and sites requiring more a tracking only solution without full third party serving. This will be covered off later in this chapter under 'Tracking but not serving'*

The assignment shown for Placement C in **Figure 3.2.4** is an absolute last resort for any trafficker to have to do; to allow them to move forward with the trafficking process they may be forced to only assign a backup image to the Placement with no corresponding Flash creative as an ad. In this instance the Placement would still require that the backup gif be supplied with a click-through URL. The result is that the user will see only the most basic backup gif. This can occur when a Creative Agency fails to supply the Flash file to the correct dimensions in time but a backup image (perhaps

from a previous campaign) is available for the trafficker to select from inside the ad server interface.

Ad Assignment

In **Figure 3.2.4** (assigned here to Placement A), it is now mandatory for just a single ad to be assigned to a placement. Indeed the flexibility is built into the ad server to allow for multiple ads to be assigned to a single placement and vice versa, saving the trafficker a lot of time:

Multiple ads to a single placement

A placement displays within it, an ad of the same dimensions but it is not always the same ad that is displayed. Refresh a Publisher page and the same placement might be displaying another ad by the same Advertiser with a different piece of creative. This is achieved through the manipulation of the controls on the placement level to change which ad is displayed inside the placement. By having multiple ads assigned to a placement the trafficker only needs to send a single tag to the Publisher at the end of the trafficking process. This might be supporting a whole heap of ads behind the scenes, controlled by the trafficker in the third party ad server interface.

A single ad to multiple placements

In some instances the same creative needs to be used across multiple Publishers but few ad servers are proficient in allowing traffickers to apply the same creative to multiple placements. Instead the ad server forces the trafficker to create a copy or duplicate ad before assigning it to subsequent placements with the same creative dimensions, even across multiple Publishers. The downside to this limitation is that any "changes" made to the creative (such as the creation

of a new creative message or the building of a new creative entirely) mean creating a whole new ad and un-assigning the old ad across multiple instances.

A single ad to a single placement
A simple sounding assignment, but one that at scale can spell trouble for the trafficker, who may have to spend hours making the assignments in accordance with the supplied media plan.

To serve the interests of time-saving, ad servers supply tools which allow the trafficker to perform these (laborious on a large scale) tasks in a short space of time across hundreds, thousands or hundreds of thousands of ad assignments (see bulk trafficking section later in this section)

A lot of ground has been covered, so here is a summary of the trafficking process so far.

a. The trafficker has logged into the third party ad server interface and navigated from the Advertiser level down to the campaign level
b. The trafficker has set up the media plan at this level by selecting the correct sites
c. The trafficker has created placements for each of the sites
d. The trafficker has found the creative supplied by the Creative Agency and created ads from them where they do not already exist as ads
e. The trafficker has added information to the ads including click-through URLs
f. The trafficker has assigned these ads with backup images to the correct placements, following the requirements of the media plan



Figure 3.2.5 *The basic trafficking workflow undertaken by the trafficker inside the interface of the third party ad server.*

Before moving onto the last step in the basic trafficking process, as shown in **Figure 3.2.4** (traffic ad tags to Publishers), the trafficker has the option of adjusting a series of the ad controls, after which it is recommended to conduct a "QA" (quality assurance) or test the placements as "tags" before they are sent to the Publisher.

Adjusting the Ad Controls

Ad Controls are essential for placements that have more than one ad assigned to them. The placement will not be displaying all of the assigned ads at once so they must be shown in turn. Ad controls are a set of logical rules that the ad server can follow when a request is received by the server to display an ad.

Frequency Capping

By limiting the number of times a user sees a particular ad the trafficker ensures that the user is not over-exposed to the same images. When combined with "ad rotation" (definition shown below) this ensures that users only see a particular creative a set number of times (called a "frequency") over a particular period of time. Frequency capping can be applied at different levels, for the purposes of this section it can be applied at the ad level but it can also be applied to the collection of ads on a single placement (called a delivery group in *DG/Mediamind*) or even across an audience (as in the *Atlas* Ad server). Although usually seen as the feature of a Publisher ad server, frequency capping on the third party ad server allows for the controlled delivery of

more expensive formats such as rich media ads, balancing format delivery to best manage advertising budget while still delivering impactful executions.

Ad Rotation

The Advertiser ad server allows the trafficker to ensure that the user is not seeing exactly the same image all of the time. By putting several images or creative into rotation on the same placement, the user has the potential to be exposed to different images and fresh content each time the same placement loads. The following describes different types of ad rotation:

Sequencing

Sequencing is the ability to storyboard a series of ads so that users that have seen the first ad in the sequence are exposed to the second in a series. Sometimes this is restricted to ads of the same dimensions and format, but ad server improvements are allowing for more flexibility in this area in the future. Sequencing uses information stored in the ad server against the user's cookie to determine which ad to show based on what ads have been delivered to them before. Sequencing is recommended for Advertiser who like to storyboard their ads, taking the user on a journey through effective creative messaging.

Weighting

Ads in rotation can be weighted so that those with a higher weighting are seen more often. Weighting is expressed as a percentage and the total weighting for all of the ads on a placement must always equal one hundred per cent. Weighting is particularly useful when optimizing a campaign to ensure better performing ads play more often.

Ad Prioritization

For some ad servers weighting can be expressed as a whole number rather than as a percentage. The lower the number, the more often it will be shown. Typically ad prioritisation exists within the band of one to one hundred. This is especially useful when ads are commonly taken out of rotation and not always replaced, if weighting was used in this situation then the new ad weights would need to be calculated each time.

Time Based

A day-part analysis of a campaign (See Chapter 5 - Advanced reports) will reveal temporal trends about best performing ads broken down hour by hour. The time based setting allows certain ads to be shown at certain times of the day. This can help significantly in avoiding "wastage" (ads being delivered to an irrelevant audience) by setting an ad to serve during quieter serving periods that perhaps costs less to serve (such as a GIF or a standard ad showing at 4am instead of a rich media ad). Also, performance may improve on different days such as at the weekend so this feature can be useful then. It is important to make a careful consideration about performance when adjusting day part delivery. Ads that fail to receive clicks at night may still be providing a hidden contribution to final sale (see Chapter 4 - Conversions).

Post Cap Gif

Some ad rotations are balanced: in which, if one ad is not shown, another is shown in the same rotation (such as with Weighting) but others like Frequency capping present situations where a user could have seen the presented ads in a finite series. Ad Servers offer the presence of an "other

bucket" in this situation (also present in Targeting) where, if the ad request is made by the Publisher page and all the ads have been shown in the series, then a backup image is shown. This image is typically a GIF and is always required in such setups. Some ad servers do not force the Advertiser or Agency to upload a single backup or default image for each placement, but instead allow them to upload a universal backup image at the Advertiser level (a non-campaign-specific piece of creative) which is just shown whenever the conditions are met.

The two levels of control

Ad controls control *only* what ad is displayed in the placement. The Publisher also has such controls in their Publisher ad server as explored in Chapter 2, but the ad ops team are using these controls to choose which advertiser's placement should be displayed. This is an complex concept but vital to the understanding of ad delivery, these are facts that often need to be re-stated during a troubleshooting process because the terminology used to describe the Media Agency traffickers ad controls and the Publisher ad ops team's controls are the same. For instance there are a minimum of two frequency caps that can be set for a placement; one in the Publisher ad server and one in the third party ad server.

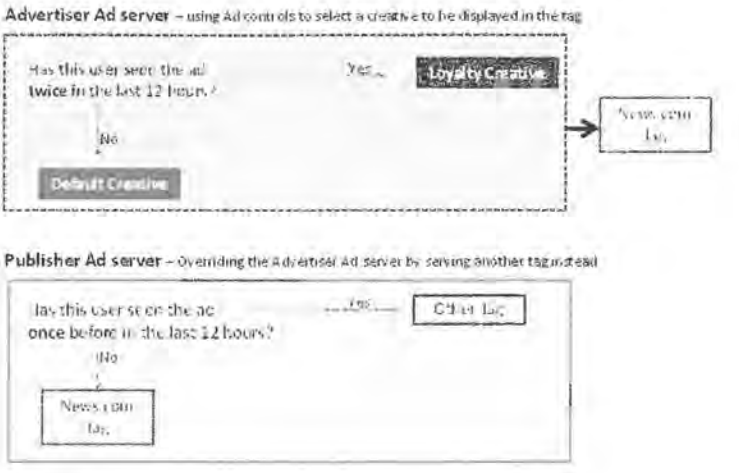


Figure 3.2.6 In the above example both the Advertiser ad server and the Publisher ad server are using frequency capping to control which ad will eventually get seen by the end user. The Advertiser Ad server can only control what goes on in the Placement (referred to in this diagram as a "tag") that it delivers to Publisher ad server. The Publisher ad server cannot change the setting inside that tag but it can choose not to serve that tag and serve another instead. In this example it doesn't matter what the trafficker on the Advertiser ad server included in the ad controls because the default creative will never get seen as long as the frequency cap in the Publisher ad server is shorter than it is on the Advertiser's side.

The ad control interface in the third party ad server is where the rules for ad delivery are managed so on top of these options is the ability to set which ad is shown to which specific type of audience.

As audiences are not a requirement for the basic setup of a display ad campaign it has been omitted from this section and added to Chapter 4 - Audiences.

After the trafficker has set the ad controls; the ad, the placement and the completed settings are ready to be used as "tags".

5. Trafficking the Tags to the Publishers

Ad Tags and Tag Formats

The creative has been turned into an ad and the ad assigned to a placement, now the "tags" can be generated and sent to the Publisher. The "ad tags" themselves are short pieces of code that simply reference the location of the ad on the third party ad server. This code is most commonly generated as JavaScript although it can also be generated as an iframe (more HTML friendly script) or for added flexibility, standard HTML.

Tags typically have a reference number located inside them which normally reflects the placement or Ad ID and are a fast track method to identify the tag for troubleshooting purposes, should problems occur later on. (The individual troubleshooting would simply take the ID and input it into a search box inside the ad server interface to locate the placement regardless of where it was located within the hierarchy).

An ad tag can be generated for each placement on the media plan within the ad server interface once the placement has been populated with ads.

```
<script src="http://bs.serving-
sys.com/BurstingPipe/adServer.bs?cn=rsb&c=28&pli=6
222360&PluID=0&w=300&h=250&ord=[timestamp]&u
cm=true&pcp=$${gmttu}$$"></script>
<noscript>
<a href="http://bs.serving-
sys.com/BurstingPipe/adServer.bs?cn=brd&FlightID=6
222360&Page=&PluID=0&Pos=2894"
target="_blank"><img src="http://bs.serving-
sys.com/BurstingPipe/adServer.bs?cn=bsr&FlightID=62
22360&Page=&PluID=0&Pos=2894&pcp=$${gmttu}$$"
border=0 width=300 height=250></a>
</noscript>
```

Figure 3.2.7 This code will render the ad inside the placement for which it has been generated when it is passed onto a web page and loaded from within a browser. It is possible to run a check on the tag on a local machine by generating the ad tag and then opening the resulting plain text file in a web browser. Both rendered image and click portions of the placement are represented in this code.

Flighting

Once ad tags have been created, the trafficker working in the third party ad server needs to get the tags to the Publisher ad server and the Publisher operations team. Few technology companies offer an integration between the Advertiser ad server and a Publisher ad server (with the most notable exception being DoubleClick) so ad tags created in the Advertiser-side technology do not automatically appear inside the Publisher technology. The ad tag code is therefore delivered to the Publisher operations team via email and the operations team receive the ad tags and put them into the Publisher ad server interface.

The third party ad server interface offers a screen at the end of the trafficking process to attach the ad tag to an email sent from the ad server directly. The trafficker inputs the contact details (email address) of the Publisher ad operations team and may attach a note, more code or other files to the email containing instructions for implementing and testing the ad tag. The instructions also contain reference information such as the placement dimensions and naming conventions, so that there is no confusion about the ad tag that is being sent.

Once the trafficker sends the tag, the third party ad server keeps a record of the send. These records are utilised by large trafficking teams on large scale display campaigns to determine which tags have already been issued. By comparing reporting data to these records later on it is possible to determine which Publishers have neglected to implement the sent tags on time.

Upon sending the tags, the Publisher ad operation team receives the ad tag to their inbox and will copy and paste the ad tag code into the technology on their side (the Publisher ad server). This will set the ad to go live in-line with the information they have received on the IO from the planner/buyer at the Media Agency or Advertiser. This information will match the line item on the Media Agencies in-house media plan. If the ad tag is not received; some ad servers offer Publishers the flexibility to log into the third party ad server directly to re-download the ad tags for their sites. In doing so the Publisher ad ops team will have third party ad server login credentials, with different permissions to the trafficker.

Thoroughly Checking an Ad is live

Once an ad tag is put into a Publisher ad server and is live on a Publisher site, traffickers are keen to check how the ads

look in the live environment on the site. Some Publishers return screenshots or send over the URL of a page back to the trafficker where the live ad can be seen. To check a correct implementation it is advised that the trafficker load the Publisher page in question while running an HTTP tracer program on their local machine. Programs like *Fiddler* or *HTTP Watch* are suitable to reveal if the correct tag is loading onto the page (by cross-referencing the Placement ID) and that both the loading of the image (or flash file) and use of the click-through URL are generating real-time server calls. Server calls which make requests referencing the placement or ad ID usually clarify that the right placement is on the page. If the wrong placement ID has been implemented onto the page the ad server reporting will not report back accurate results and the Publisher ad ops team should fix the tag on the page with the correct placement ID.

In addition to running these manual checks, the trafficker can use the third party ad server to verify actual delivery to the right user audience such as users in the correct geography. This is called Verification and further information about this and the monitoring of delivery information is available in the third party ad server reporting interface (see Chapter 5).

Once the correct tags are live on the Publisher pages the campaign setup is complete and the trafficker will access the reporting interface on the third party ad server to begin the process of Optimization (see Chapter 6).

7. How does a trafficker manage campaigns with hundreds of Ads?

Bulk trafficking through *Excel* templates

With thousands of ads, placements, sites and many assignments, un-assignments and amendments to be made, bulk trafficking tools are indispensable for the trafficker. The trafficker can download the ad server's trafficking sheet from the media plan level to a local computer and fill it in row by row before uploading it back to the ad server via an error checking and queuing system. The trafficking sheet is designed so that it can be read directly by a machine. At any time the full media plan can be exported from the ad server interface and appended before being re-uploaded so that bulk changes can be made. The ad server makes a log of changes at scale by specific traffickers so that if multiple traffickers are working on the same media plan at any time, there is little confusion about which trafficker made the last changes to the media plan.

Bulk trafficking with *Excel* as the UI

Some tools like the *Atlas* 'Media Add-in' mean the trafficker can complete almost every task of the trafficking process from inside Excel including exploring the ad server hierarchy. Like the Flash components supplied by the ad server to the Creative Agency developers; an additional piece of software can be downloaded from the ad server and installed locally on the trafficker's machine to enable the use of the bulk software. Essentially a skin or plugin accessed directly from the excel interface, permits the execution of various operations when trafficker credentials are verified. The gain in operational scaling is very apparent but the downside is that the interface of an excel spread sheet is a tenuous replacement for the usable designs of ad serving interfaces. As of publication, the marriage of a true bulk trafficking solution and an ad serving interface has yet to announce

their effective engagement beyond a spread sheet-style design.

Automated trafficking setup via ad server API
It may seem strange that it is possible to transform Excel into a user interface for the ad server, but as long as a connection is established with the correct user credentials, excel and other programs can be used to replicate the actual ad serving interface. This is made possible due to the ad server "trafficking API".

An API (or Application Programming Interface) allows an ad, web page or program at a distance (external to the server in question) to request only a select piece of information from the server and have that information returned in order to fulfill a function.

Services that offer an API quickly ensure the propagation of their information around the internet and soon become indispensable to the workings of the web. Take for instance *Facebook* which allows developers to sign up to their API to retrieve information about users (with their consent). Or *Google Maps* which allows anyone who has requested a free developer 'key' (a type of password) to display a *Google* map based on the longitude and latitude information supplied, passing the developer key in a request to the server directly.

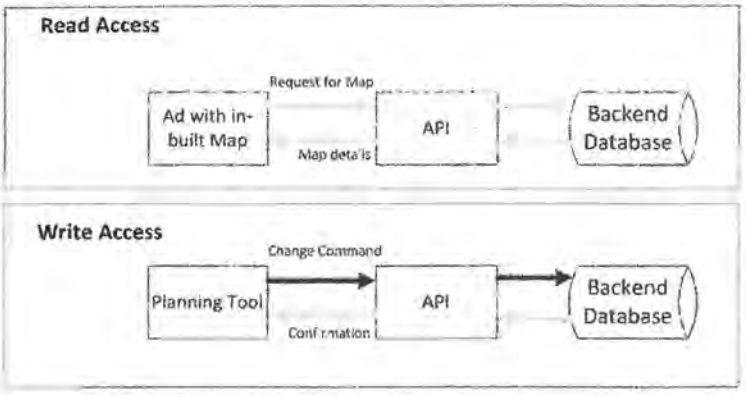


Figure 3.2.7.1 With just read access an Ad or any other application can make a request to the backend database via the gateway of the API. The Request for Map is joined by the API key which grants access to the backend database. The backend database processes the request and if the API command is correct and the key is accepted, it returns the required details which in this instance are the details of what the map looks like which are subsequently shown in the ad. With write access an application such as a Media Agency planning tool can issue a command to the ad server database via the API such as to have a new placement trafficked. If the API key and the command are accepted then the database makes the amendment to the setup of the ad campaign and returns a confirmation to the planning tool that the task has been completed.

As **Figure 3.2.7.1** shows, with an ad server trafficking API, the external program can issue commands to the ad server from a distance to request that an action be executed or for the submission of information into a particular section of the hierarchy. This two-way functionality of both affecting setup in the ad server media plan and returning confirmation of the changes is a step further than the "read-only" access of

public APIs offered by *Google* and *Facebook* in the previous examples.

The ad server trafficking API supports a multitude of commands and can return useful pieces of information about the campaign setup. This is how it is possible for a remote copy of excel to make requests to the ad server and fetch a current copy of the media plan without requiring the trafficker to leave the trafficking interface of the excel environment. In the future, strong ad server APIs supporting all functions of the ad server will permit application developers to create simplified user interfaces and dashboards which could be installed on mobile devices allowing traffickers to traffic, monitor and control campaigns on the move. For the time being trafficking APIs are predominantly utilized by in-Agency planning tools and tools like *iDesk* (as touched on in Chapter 2 - Agency Planning).The media plan created in *iDesk* makes the relevant calls to the ad server using it's API and entities such as the placements can be created from afar. In turn the ad server can return information such as a list of available "site" entities so that *iDesk* can select one and return that request back to the ad server: trafficking at a distance.

Other API capabilities exist to support the ad server capabilities such as for reporting. More details can be found in Chapter 5 - Reporting.

The JavaScript carriage capability

Another of the more common advanced capabilities of the standard display ad is its capability to launch additional code on the condition of certain JavaScript based activities resulting from real-time user action. Such activity may include the JavaScript parameter of "Unload"; whereby on the condition of the ad being abandoned (such as by leaving the page or closing the browser) the JavaScript can in turn

launch an on-screen survey. This particular example is the delivery of what is referred to as "Survey Vendor code", but there are many actions that can be taken, allowing the trafficker to manipulate the JavaScript capabilities from the comfort of the ad server interface.

Survey vendor code and third party scripts-in ads

Survey Vendors emerged in the world of ad serving technology because true Advertising effectiveness is a subjective concept. Although an ad may have performed well according to ad server reporting (receiving clicks, interactions and even conversions (see Chapter 4) it is still not one hundred per cent clear if users were influenced by the ad itself to perform these actions. It is entirely possible that if the ad was shown in the right place at the right time, that that ad merely acts as a convenient pathway for the user to complete an action they had every intention of doing anyway.

Furthermore it is not clear if the message inside the creative itself had an impact on the user. Since we cannot see what is going on in a user's head, collecting this information requires some qualitative surveys to be conducted.

Research companies or Survey Vendors such as *Dynamic Logic* create short surveys that can be launched as part of the delivery of the ad, to collect this information from users. When the ad loads, is clicked or closed a small survey loads onto the screen. This survey overlays the rest of the information on the Publisher page and is designed to quickly ask the user about ad effectiveness, without disrupting the user experience.

This information, once submitted in the online-form, is sent on to the Survey Vendor's database on an external server. In order for the survey to load up in the first place, a small

piece of code needs to be provided by the survey vendor to the trafficker long before the campaign goes live. The trafficker then loads this code into the third party ad server interface at the ad level when trafficking the campaign (adding the code to which ads it is required to fire off against). Pre-created fields in the ad server interface allow for these codes to be fired during one or more of the JavaScript activities on the placement.

To setup a Survey from the outset of the campaign, the Agency/Advertiser needs to reach out to the survey vendor when the campaign is being devised and the supplied code by the vendor should deliver some unique identifier in its return ad-call that will allow the data from the vendor to be tied to the data from the ad server later on. A Placement ID is suitable for this purpose. Therefore an optimal setup process would be:

a. The Placements should be created by the trafficker first
b. Then the trafficker should export the media plan in the bulk trafficking tool (displaying all of the Placement IDs) and send it along to the survey vendor
c. The code corresponding to each of the placements should be created in the Survey Vendor interface (not found inside the ad server), with a matching Placement ID
d. The survey vendor should update the "survey vendor" field in the supplied exported media plan before sending it back the trafficker
e. The trafficker should ensure that the code is set to fire at the right time (by either adding the information in the media plan and then uploading the media plan back into the ad server interface, or by making the necessary adjustments on each placement once the media plan has been uploaded).

The various nuances of the trafficking workflow are designed to facilitate clever additions to the process, such as the addition of survey vendor code. The availability of the bulk media plan functionality makes it much easier for the trafficker to work with these additional trafficking requirements.

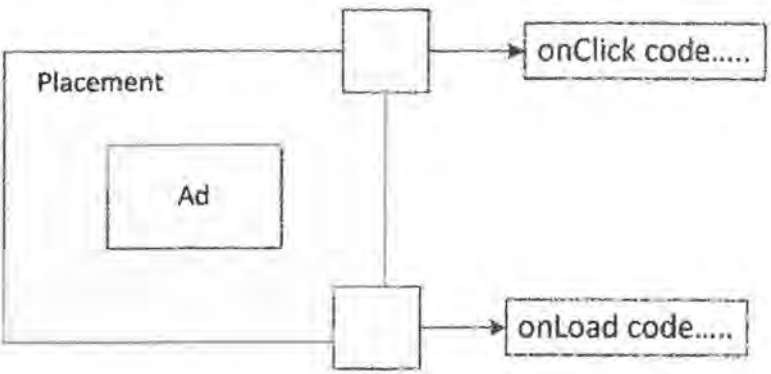


Figure 3.2.8 Placements and Ads all have the capability to support or provide a "piggyback" to other blocks of code. These codes can be added by the trafficker during the trafficking process and appear to the end user as seamless functionality to the ad, if they are seen at all. These methods allow ads to act as channels which can pour information to more external sources. Publishers often keep a keen eye on ads for this exact purpose because ads pouring information to other networks can tarnish the Publisher's own valuable audience - a process known darkly industry-side as "data leakage" (see Glossary).

We have touched on the ability of the ad or the placement to support the "piggybacking" of additional code (albeit the survey vendor code for now). The same fields in the ad server interface which allows a JavaScript activity such as "onLoad" to be selected on the ad level, before firing additional code, exist at the ad level, placement level and

higher up levels such as the campaign level. The campaign level allows the same code to be applied to all ads in the campaign. The Advertiser level allows the code to be applied to all ads in all campaigns for the whole Advertiser. If the trafficker pastes code in from a different system, like Survey Vendor code (called "third party code" or "4th party code") into these higher levels of the hierarchy, their use allows for the bulk implementation of such code, feeding the fields in the entities that trail below. This paves the way for a concept touched on in Chapter 1 - Double ad tracking.

Recalling from Chapter 1; double ad tracking means that a single ad server can have a view over activity running in more than one ad server. This is made possible because the ad tag code from one ad server is exported and implemented by the trafficker into the same third party code fields that would be used for survey vendor code.

It would not make sense to have two whole pieces of creative from a both ad servers appear in the same ad spot at the same time, so the trafficker in the lesser ad server needs only to supply the most basic type of tracking. The creative will then just be served from the one system, but the reporting of the activity will appear in both. In this sense an ad server's primary function is to serve the ad but in a close second is the function of tracking. This same method can be used when a Publisher decides to serve the creative from their own system.

Tracking but not serving

It is quite common in display ad serving to have a fragmented media plan; where the trafficker has creative for most of the setup but must use "tracking pixels" for the rest. The creative is usually served from the third party ad server directly, but some Publishers 'do not support third party ad serving' and instead permitting the Media Agency trafficker

to supply a tracking pixel (for each of the Publisher-hosted ads) so that the trafficker can be sure that the activity is tracked or represented in the third party ad server.

Is it possible that a Publisher cannot support third party ad serving?
Publishers have less to gain from the use of a third party ad server in a campaign than Advertisers or Agencies do. Implementation of code requires operational resources, extra training and ideally the presence of a Publisher ad server. These may not be available to small scale Publishers or Publishers with content housed in very unique environments (such as the emergence of ads in mobile apps created without a standardised SDK - see Chapter 3 - Mobile Ad serving for more information). Clashes between code supplied from an external system could in theory also conflict with the Publisher page, particularly if there is a poor CMS in place. However given the demands of Advertisers to utilise third party systems to verify delivery and performance, poor systems in place on the Publisher side are bad excuses for not offering total transparency. Publishers which refuse the use of third party ad serving tags are risking deeper investigation by the wider community to uncover reasoning to avoid transparency.

Permanent tracking pixels need to be created in the ad server by the trafficker and their ad tags sent and inserted into creative (served by the Publisher directly by the Publisher ad ops team, from the Publisher ad server). This brings the counting methodology as close as possible to the third party method. In this instance the trafficker can still send the ad tags to the Publisher at the end of the trafficking process, but since it is not strictly displaying the ad, the code looks a little different. Some ad servers refer to this code

package as a "Tracking Placement" and it is made up of a tiny 1x1 tracking pixel image and a single line click through URL (a click redirect). The trafficker will need to create one for each corresponding item on the Media plan, where the ads or assets are being served by any system other than the third party ad server. To keep things nice and tidy, the tracking placement should be treated like any other placement in the platform and be assigned to the correct site. Given the nature of tracking placements it doesn't make sense to add the "ad control" options to these entities.

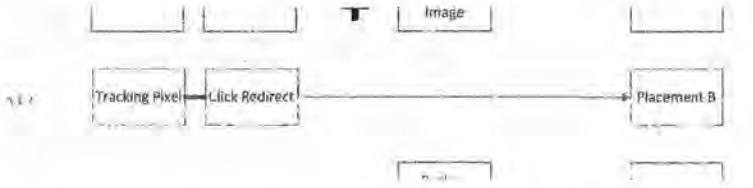


Figure 3.2.9 another ad assignment option for the trafficker: assign a tracking placement (a tracking pixel with click redirect) to a placement. The flexibility of a tracking placement allows the tag to be trafficked to a second system where the actual ad or creative might be hosted. When the tag for the tracking placement is trafficked it is consider to be a "4th party pixel" in any third party system.

A note on Click Commands

DFA (the *DoubleClick* third party ad server) allows traffickers to create multiple "click commands" (Click Commands are just click re-redirects) but crucially allows the assignment of many of these entities to a single placement. When these tags are trafficked to the Publisher, the Publisher has the freedom to choose just one of these click commands from the supplied tag (where the Publisher may be site-serving the actual ad). As each of the click commands might point to a different landing page the Publisher can use the power of the Publisher ad server to rotate the click commands across

the site-served creative for the purposes of optimization and landing page testing.

Since all of the click commands are assigned to the same placement, *DFA* records these "click" metrics under the same placement in ad server reporting. This gives the Publisher more freedom in determining which ad to display while site serving while still counting clicks in the correct place in the *DFA* platform. Ad servers which limit the creation of one click command to one tracking pixel in a tracking placement format do not force the Publisher to use both image (pixel) and click portions of the tag. The tag will work and the metrics will count even if only the click portion or just the pixel portion is used.

With the combination of these more advanced tactics in the trafficker's arsenal, most display campaigns can be setup and managed. Before closing out this section let's focus briefly on the mechanisms at work here which are useful in determining why an ad that may have been trafficked already and implemented onto a Publisher site is either not reporting or displaying as expected.

7. Ad Replication

Ad Delivery and Replication

To get an ad to deliver all over the world requires several technologies working together. The ad server itself will normally do the complex calculations and decision making about which ad to display as well as storing the images and flash ready for distribution, but getting the ads from the ad server and its hardware to billions of pages (and eyeballs) requires the support of a CDN or content delivery network (See Chapter 1).

The trafficker has the correct access to the ad server interface to swap in and out new ads even when the campaign is live but this action does not have an immediate result. The time it takes to swap a creative over (or replacing an old ad with a new one) and have it appear all over the world in a live campaign is referred to as "ad replication", and for far reaching global distribution campaigns should take no longer than around fifteen minutes (after saving the changes in the campaign management screens of the ad server).

For those looking to check that such changes are now "live" it is possible to go to a Publisher page and refresh it until they see the change live, but it is recommended to clear browser cache and cookies (and waiting fifteen minutes) before refreshing the page. This ensures the old ad has not been cached locally and removes the restrictions placed by frequency capping at the Advertiser or Publisher ad server.

Finally if the ad is still not displayed, it is wise to check the settings for IP (geographical) targeting since a tester sitting in the wrong geography will not see the correct creative if the campaign is working as expected (more on this in Chapter 4 - Audiences).

Server-Side Caching

In addition to this there is a caching mechanism on the ad server itself, meaning that if a tag that looks the same is served frequently, the ad server will make and store its own temporary copy for faster delivery. To eradicate the temporary copy in the ad server, most tags include a string or token in the pixel portion of the tag code itself. This string is simply a random number that changes every time the ad is served out, and since the random number is never repeated, the cache on the ad server is eradicated. The string is normally represented with something that looks like:

"***?rand=RANDOMNUMBER***"

When the trafficker is piggybacking a 4th party tag, the instructions in the third party ad server will explain that RANDOMNUMBER in capital letters can be over-written with a system-friendly macro which, when the pixel is fired, will insert a random number in automatically. Inside *DG Mediamind* this macro is: %RAND%

What exactly is a tag?

When the tag code is sent onto the Publisher the implemented tags are conduits for information being sent to the ad server for reporting and creative delivery purposes but the tag itself is a very small a light piece of code and stand-alone it does not contain all the code it needs to perform certain functions. As such, the tag calls upon one or more JavaScript libraries which contain more complex commands (and are a touch heavier in terms of file size). The construct of these libraries is unique to each ad server and some make better use of various coding methods than others to ensure a slicker and smoother experience with say; video ads, for instance. Understanding how these parts work, detracts from the more immediate focus of this text but third party ad serving providers will be able to supply this information. The important thing to note is that more complex capabilities are only possible through the enabling of these libraries, therefore any tags implemented in a more simple tag format such as a simple tracking placement as opposed to a fully-fledged JavaScript tag will have certain functions missing (including viewability).

9. Making a bigger impact with Display Campaigns

The best advice to improve knowledge and skills further in standard display is to undertake some trafficking in the ad serving interface itself. Most third party ad servers provide a wealth of training materials and support to guide the trafficker through the campaign management process in their specific system and will be happy to assist in the setup of a few test campaigns for training purposes. The good news is that all the concepts covered are common to all the big third party ad servers, those absent features in one system will soon crop in the others, play a game of catch-up with one another to provide more and more standardised features.

The next section will look at the Rich Media channel. Rich Media is an advanced form of display advertising, designed to engage and attract the user. It is especially popular with big brand Advertisers and is the gateway to Video ad serving. Due to its powerful impact Rich Media is likely to see a "come-back" in the next few years, emerging in innovative ways through video players, mobile and through more advanced forms of reality augmentation, but its technical setup is in some ways a little more challenging than the setup for standard display ads.

Chapter 3 - Section 3
Rich Media

- What is Rich Media?
- Why use Rich Media?
- Core Formats
    - Expandable ads
    - *Figure 3.3.1a*
    - Floating ads
    - *Figure 3.3.1b*
    - Video ads
    - *Figure 3.3.1c*
    - Pushdown and Sidekick
    - *Figure 3.3.1d*
    - Data Capture
    - Home Page Takeovers, Wallpaper and Page Skins
    - *Figure 3.3.1e*
    - *Figure 3.3.1f*
    - Polling
    - State Based ads
    - Local Connection ads or Sync Ads
    - *Figure 3.3.1g*
    - Print ads
- Standardisation and Publisher certification
- From the top: Building Rich Media ads
    - Training the creative developer
    - Two schools of Rich Media specs: Publisher and Ad server
    - Parent and Child files - the polite load
    - Expansions and Floating -the declaration of state
    - Interaction tracking: Default and Custom Events
    - Interactions as additional clicks
    - Building Custom Rich Media ads
- Qaing and Testing Rich Media ads
- Trafficking Rich Media ads

- Publisher Troubleshooting
- Rich Media for Mobile & Rich Media for Instream

The following section looks in depth at the full campaign setup process for Rich Media ads which are served and tracked by the third party ad server and trafficked to Publisher websites.

Strictly speaking Rich Media is not a channel but a subset or evolution of Standard Display ads. This means that a large margin of the trafficking process covered in the section on Standard Display also applies to Rich Media. With this new territory comes new terminology and concepts that begin right at the moment of creative conception and trickle down to the delivery of the ad in the user's browser.

What is Rich Media?
Technically Rich Media refers to an ad that contains one or more of the following features:

- Video
- Expandable panels
- Interaction Tracking
- Polite Loading
- Floating Effect*

(* Floating ads are sometimes referred to as "Out of Banner "ads, although this phrase has also been coined to describe all types of Rich Media ads in some ad servers)

Why use Rich Media?
As outlined at the tail end of the last section, Rich Media offers Advertisers the chance to make more of an impact with creative that is more engaging for the user than standard Display ads or Search ads. Rich Media is really collaboration between the third party ad server and the Publisher. Implanting a Rich Media ad onto a Publisher website is going to draw user attention towards the ad due to the nature of its animations, in doing so it is

possible for ads to be created that completely hide the Publisher's content. This is where the phrase "out of banner" comes from as some Rich Media ads literally spill out of the space they were designed to fit inside.

Since Rich Media causes a longer loading time on a page and potentially could conflict with the Publisher's own content and code, the Publisher must be on-board to accept Rich Media before it is inserted into the Media plan for an Advertiser campaign. Typically a Publisher will offer Rich Media at a higher CPM to Advertisers and Agencies for specific ad spots. Ads occupying these spots will need to adhere to set of specifications defined by the Publisher which will have been created to ensure that within the confines of the specifications, the page and the user will not be overtly affected by the engagement with these ads.

At a higher cost of CPM and a bigger drain on the marketing budget for the campaign, one might question how to prove that Rich Media is offering more value than standard display. The answer comes from the possibility of tracking interactions. As detailed in Chapter 5 - Reporting, Rich Media goes beyond impressions, clicks and conversions to provide metrics that explain both what the user is engaging with and for how long. These triggers and timers allow the Advertiser to determine how a user engages with the creative (allowing for improved creative optimization) and provides a new window on effectiveness by looking at "engagement" on interaction as opposed to clicks to assess performance. It is safe to say that when looking into the working parts of a Rich Media ad, there is more complexity than there is for a standard ad which means a greater opportunity for creativity. Greater creative paves the way for awards for creative developers and designers which in turn attract good PR for Advertisers and for Creative Agencies. The most impressive designs eventually become adopted at an industry level and it is thanks to the standardisation of formats by the IAB and

collaborating publishers that we now have formats such as the "Rising Star formats" which circumvent the standardisation of Rich Media from becoming too diverse and unmanageable from both a development and testing perspective.

Core Formats

Formats are the core of what makes Rich media so scalable. By focusing on the features of the most common rich media formats it is possible to identify what kind of format a Rich Media ad is just by looking at it or playing with some of the controls. Formats are really just descriptions of advanced functionality within a Rich Media ad but the presence of established format names moves allows planners, buyers, marketers and traffickers to have meaningful conversations about rich media campaigns with little or no knowledge about how the ad works.

With an established set of format names, it is possible to use the naming convention used by the creative developer in the creative file name, or the name on the media plan or Publisher IO, to ensure that all parties have a mutual understanding about the formats used in a campaign. Formats are not universally priced, more advanced functionality typically means a higher end cost to the Advertiser.

The following is a set of the most common formats found in Rich Media:

Expandable ads

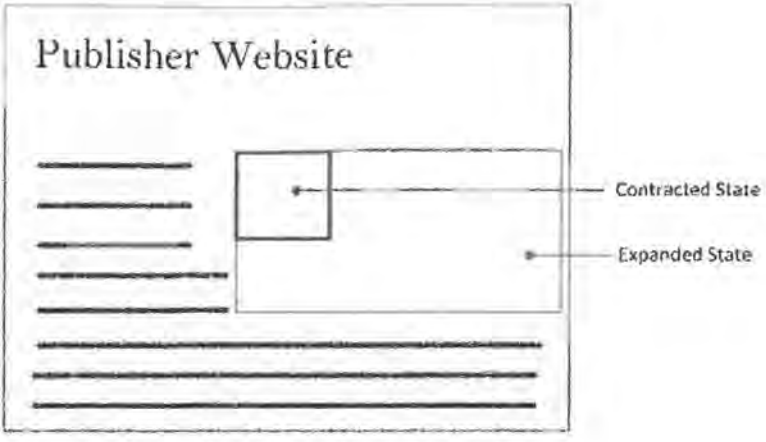


Figure 3.3.1a *an expandable ad grows in size over the publisher content, overlapping it briefly while the ad content is delivered to the user. This figure depicts Rich Media Ad states of expansion and contraction which are covered in more detail later in this section.*

One of the most commonly used formats in Rich Media is the expandable (hopefully in time a word that will find itself a place in the Oxford English dictionary following ten plus years of industry acceptance and use). This is a format that grows in size ("the expanded state") and then contracts ("the contracted state") down to a fixed ad position (called the "ad location") on the page. Expandables can "auto expand" (if Publisher specifications approve) meaning that they start playing in the expanded state rather than requiring a user action to cause the expansion. An expansion can also be triggered by a click or mouse rollover. They can also expand in any of the four directions or expand in all directions at once. Expandables can have a transparent background so that in the expanded state the animation appears to interact with the content of the page itself.

Floating ads

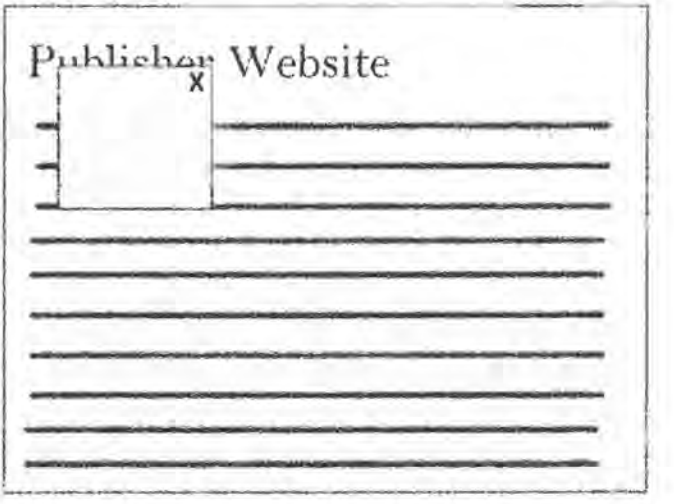


Figure 3.3.1b *The floating ad format loads over the Publisher page content, sometimes tinting the page behind it grey until it is closed with the X shown in the corner. Once closed, floating ads do not strictly return to any fixed point on the page.*

A diminishing format in both the UK and US markets; the floating ad appears on top of the page content, not tied or anchored to a fixed ad spot. The ad itself does not open in a new browser window but often comes with the same qualities as a new browser window (it can be dragged around and closed). The floating ad goes through various states of play during which a user can choose to close the ad via an X in the corner. Coming across a floating ad as a user without an X denotes a failure on both the part of the creative developer and the ad server for failing to notice its absence (this helps explain why the Rich Media QA process is so important). Floating ads are typically frequency capped at the Publisher ad server level, so that if the user loads the same page several times they will eventually stop seeing the

floating ad so as to sustain a reasonable user experience and not risk "creative fatigue".

Video ads

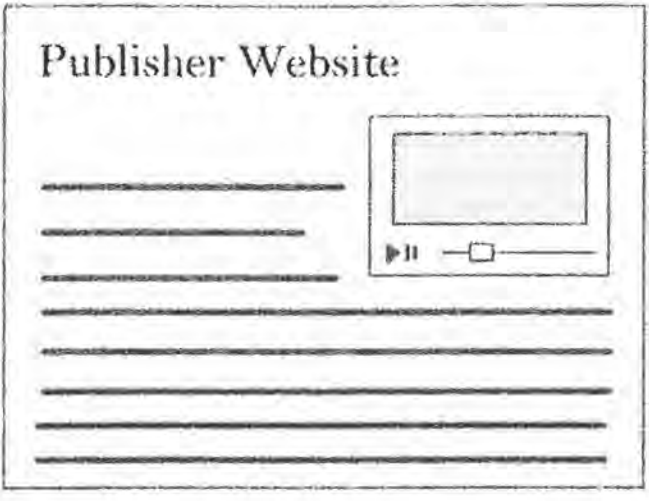


Figure 3.3.1c *Video ads appear in a standard ad spot on a Publisher web page typically playing a short video clip with controls for the user. Most commonly used by film studios to display trailers.*

Video ads are one of the most commonly used Rich Media formats and can be combined with the expandable or floating features. Video ads look like standard ad units but play video (usually with some controls) from inside the placement. Video ads should not be confused with Instream ads (which are typically traditional TV ads purposed for online video broadcasters and purpose built video players - See Chapter 3 - Instream and Video). In Rich Media Video, videos can start automatically but the muting functionality is dictated by the Publisher specifications for which the requirement is most common on mute. HD video, green screen video, full screen video and multiple video play (where many videos appear in the ad at once - like in the

case of the format called "the video cube") are all supported by most third party ad servers. Video can also be streamed into the ad live, and the stream itself can be loaded in real time or pre-loaded into the ad by a method referred to as "progressive streaming". Specialist technologies and players have emerged to support unique Advertisers needs such as live video feeds: one such example is *Brightroll*.

Pushdown and Sidekick

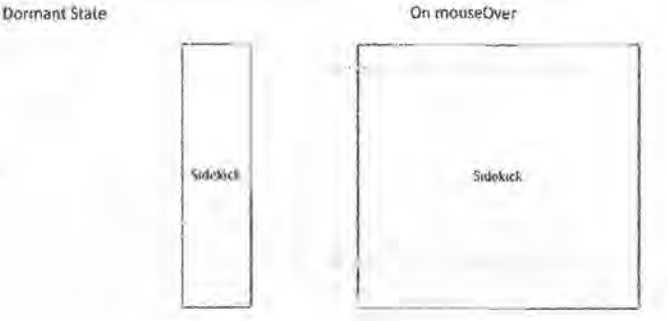


Figure 3.3.1d *A Sidekick format before and after the user rolls over the ad with the mouse. The rollover causes the side kick to expand, seemingly pushing the content of the Publisher page to one side rather than rolling over it as it would normally do in a standard rich media expansion.*

Innovative Rich Media formats are constantly being developed but finding ones that are engaging and not obstructive to the user experience can be a challenge. The pushdown and sidekick IAB approved formats allow the ad to physically push the page content down or to the side in a hybrid of the traditional rich media expandable ad. The user is exposed to the larger canvas of the ad temporarily until they request the content be returned by a button or slider. Such formats are even more effective on devices without a mouse where a sweeping action on a touch screen is more

intrinsic to the user experience with the ad. Ad Serving providers and specialist Rich Media technologies such as *DG Mediamind* and US-centric *Pointroll* work with major Publishers to develop new formats with the aim of better monetising the Publishers inventory. Formats chosen by the Publisher can often be made with a contract of exclusivity with the technology company that developed the format for a fixed period. In such instances any Advertiser or Agency running their ad in a new and exclusive format might have to ad serve their ads through the advised technology to have it accepted.

Data capture

Owing to their larger file size, Rich Media ads can do more than standard flash ads, even going so far as to replicate the full Advertiser website in the confines of an ad. A good example is with a Data Capture form which passes or collects data about a user via a form in the ad and sends it off to an Advertiser database in real time, potentially returning information like a voucher code. A use case would be for a mobile provider like the European mobile giant *Telefonica* offering free sim cards to users.

Since no payment details are required, the user completes a form in the ad and submits it before they are given a confirmation message in the ad. The confirmation message contains a conversion tag (see Chapter 4) which logs a conversion without needing to direct the user to the Advertiser site and by returning the confirmation, the data capture form has spoken in real time with the *Orange* database clarifying that the free items will be sent to the address details supplied. Data capture forms of this nature are not meant to store Personally Identifiable Information (PII) on the ad server, but instead encrypt the supplied information and send it to the Advertiser database directly in line with local data protection laws.

Home Page Takeovers, Wallpaper and Page Skins



Figure 3.3.1e *a wallpaper execution sits behind the Publisher web page stretching from the left to the right gutter. When further Rich Media functionality is added to the default execution the effect is engrossing including animations which appear to pull the whole of the Publisher page apart in its entirety. Less sophisticated versions simply place very large images in the background. Typically the Publisher page is not transparent as shown in this diagram.*

Modern Publisher site designs centre the content on the users screen producing large "gutters" of empty space down the sides. These can be filled with a giant image, images or a playing video (a very nice effect on a big screen). This format is referred to as a Wallpaper. In comparison a Page Skin tends to mean advertising in every possible ad space and even sponsoring Publisher content working in close collaboration long before the campaign is due to begin (sponsored content is not yet served by third party ad servers but is referred to as an "Advertorial"). The presence of a Page skin is also referred to as "skinning" a page although when it comes to skinning a homepage, it is

referred to as a "home page takeover" (HPTO). Home page takeovers are a big deal - they are usually more expensive than any other formats because of the enormous reach they have (think *MSN* or *Yahoo!* Homepage) and rarely last longer than a single day. HPTO's alongside Instream ad units are the cream of what is referred to as premium inventory (the most valuable ad space money can buy).

HPTO's bring a wealth of requirements and needs to the third party ad serving technology (such as real time reporting and optimization, as well as the need for the ad server to have strong relationships with big Publishers to ensure that technologies work in harmony with one another, even when Publisher pages are re-designed).

HPTO's are not as simple to formulate as they sound. For instance the Homepages of a major Publisher like Yahoo! differ in each country, therefore the HPTO format for an international campaign, may have to undergo an extended period of development to suit the homepage of each country.

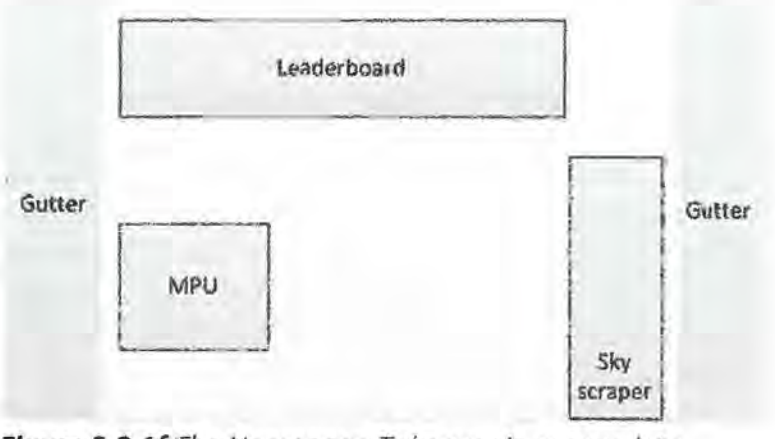


Figure 3.3.1f *The Homepage Takeover is a complete unification of all of the available ad spaces on a single Publisher web page used in a single campaign by a single*

Advertiser. Homepage takeovers are usually high cost executions that require strong collaboration between the ad server, the Publisher and the Creative development team.

Polling

A poll ad can take user information entered into a single ad and display the results of that collective input in real time inside the ad to all other users. Behind the scenes, the ad server is often powering a small database which collects the results purely for the poll.

This collaboration effect combined with a high reach campaign, can achieve poll results that are completely unattainable by a poll on an Advertiser site alone. Some clever manipulation of this technology can connect thousands of users together in real-time collaboration. So instead of collecting simple survey results from individual users and displaying a graph to the collective, users can take webcam photos of themselves and collectively build a collage.

State-based ads

Not strictly a format but rather a feature of Rich Media ads is the concept of "State-based ads" which are a more flexible way of performing creative sequencing than using the in ad-server sequencing options alone (see Chapter 3 - Display ad serving). State based ads simply remember what "state" the user last put the ad into (such as leaving an ad halfway through the playing of a video) so that should they encounter the ad or campaign or even Advertiser again, the next ad which follows could pick up from where the user left off or perform another action.

State based ads work by dropping an additional cookie on the user machine so that the "state" can be recalled later or

utilise a reserved area of local memory available for which both Flash (Flash Cookies) and HTML5 can take advantage.

Local Connection ads or Sync Ads

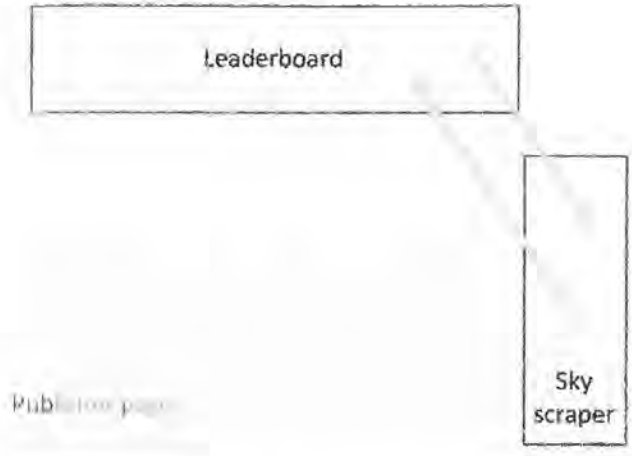


Figure 3.3.1g *Sync ads appear to communicate with each other on the page creating an affect that the two separate units are really just one ad. A clever manipulation of the Flash Actionscript code combined with the capabilities of the ad server allows the ads to appear to be completely synchronised.*

Local connection ads, also called "Sync ads", create the wonderful effect of the leader-board or masthead and the skyscraper unit speaking to each other in tandem. The surprise dawns upon the user when the seemingly separate ad units both displayed on the Publisher page at the same time, suddenly interact with one another such as passing an object between them or syncing in time with each other. For the Publisher ad operations team it is just a case of implementing two separate ad tags on the page at the same time. Rich Media templates created by the third party ad server allow the Creative Agency to swiftly and easily

reproduce this effect for those Publishers that support a matching layout. Sync Ads are more commonly used in the process of skinning a whole page, since two of the major ad units are already taken care of in a unique way.

Print-ads

As explained later in this section rich media ad units can contain a number of hidden away files as "child files" which are initially displayed to the user. One of these so called "child files" could be a full size print version of a retail voucher, which inside the confines of a small or even standard rich media unit has trouble being displayed for printing purposes on a small screen. The rich media unit can be coded to allow the user to take advantage of the voucher offer from inside the rich media creative and hit a print button, to print out a copy of the hidden child file and thus the voucher. Again this is more of a Rich Media feature but gives a good indication as to what else can be stashed away utilising the much larger file sizes granted with the use of Rich Media capabilities and specifications. Whole brochures in PDF format are not uncommonly used to stash away more information for the user to expose later on, inside a child file.

A run through of popular formats is rather like browsing a menu in a restaurant. Different dishes suit different Advertisers at different times. There is no hard and fast rule for what to use for the right impact or message. Often the skill of the creative developer or agency has the biggest influence in this area. It is all well and good knowing what is available but it is important to ask two key questions before moving forward with a Rich Media campaign:

a. Can my chosen format be run on any site or page? (Specifications and Certification)

b. How are these culinary delights created? (How are Rich Media ads built?)

These considerations will be broken down in sequence in the following paragraphs:

Standardisation and Publisher certification
In the early days of Rich Media, even with the *IAB* trying to standardise formats, many high volume traffic sites on the internet had just not encountered Rich Media ads in the past and although they were keen to run new formats, these had to be tested with test creative and tags before the campaign could go live. Rich Media technology providers and third party ad serving providers employed teams of people to contact Publisher sites well ahead of a campaign launch date on their very first engagement regarding a specific format to test that it would work on the Publisher site before the campaign went live. With rigorous testing the formats could be run and the same format would not need to be re-tested for future Advertisers and campaigns. The Publisher site was therefore "certified" to run the format.

Today big Publishers have gotten far more efficient at this process and Rich Media certification for formats deemed as standardised by the *IAB* no longer require certification in this way. What this means it that as long as Publishers and technology providers (including third party ad servers) serve the *IAB* Rich Media formats in the way laid out by the *IAB*, the standardised method means that no further testing is needed.

This industry development made Rich Media scalable and the process of standardisation and certification means that any Publisher subscribing to the standardised practices of the *IAB* can run *IAB* Rich Media formats. The same goes for the ad server. The practice of certification still goes on where standardisation and adoption has not reached full industry backing because the channel in question is too new: today this is true of certain

formats for Instream and certain formats of Mobile (see Chapter 3 - Instream and Chapter 3 - Mobile). These newer formats are the living evolution of traditional Rich Media so it makes sense that the industry will return to the certification processes that worked so well in the past.

From the top: Building Rich Media ads

The building of Rich Media ads follows the same development cycle as building standard display ads generally speaking. The differences are building in and testing the more complex elements of the functionality. The final part of this section will cover off the additional considerations that need to be made when building Rich Media ads.

1. Training the creative developer
2. Two schools of Rich Media specs: Publisher and Ad server
3. Parent and Child files - the polite load
4. Expansions and Floating -the declaration of state
5. Interaction tracking: Default and Custom Events
6. Interactions as additional clicks
7. Building Custom Rich Media ads

1. Training the creative developer

Creative developers who have built ads for ad servers within the Flash authoring environment in the past will still need to be introduced to the realm of Rich Media ad creation since the various Flash elements (Expansions, Floating, Video, Interactions and Polite loading) utilise different aspects of the components supplied to the creative development team by the ad serving technology. When it is first clear that Rich Media will be used in a media plan, the technology provider hosts what is referred to as a "kick off call" between the Media Agency and the Creative team. This meeting is used to allow the Media Agency to explain to the Creative Agency

what execution or format had been agreed to run with a given Publisher and the creative developer is subsequently asked if they need additional training on the available tools in order to build the required format. Creative developers should usually be trained an absolute minimum of a week before the ad needs to be completed, to allow the developer time following the training to get familiar with the new aspects of the supplied components and to work out how to get creative with the code. The kick off call is used to allow the Media Agency to explain what their expectations are from the execution in terms of ad server reporting during and after the campaign. In some instances this will require the developer to undertake additional work implementing what are referred to as "custom interactions" or "custom events". Before the work can start however the creative team always needs to know one thing above all else: what are the limitations to the creative? This is vital since Rich Media opens the doorway to more advanced creative executions, so it is up to the Media Agency and third party ad serving provider to lay down the ground rules or more advanced specifications from the outset.

2. Two schools of Rich Media specs: Publisher and Ad server

Ad Server Rich Media Specifications

Creative developers and designers approaching rich media commonly ask for the ad server "specs". Essentially the request is to determine what needs to be done to the Flash file creative so that is accepted by both the third party ad server and the Publisher. Third party ad servers and rich media technology providers make the work of a developer easier by supplying a set of Flash Components which can be installed in the developers local version of Flash (just as for Standard Display ad serving). The components along with a supply of templates, allows the developer to "build" the ad, containing the code required to upload the creative to the

ad server. This code is the majority of the "specification". Casting one's mind back to Standard Display, the only code required by most ad servers is the presence of the "clickTag". With Rich Media this is often such a lengthy chunk of code siphoned into multiple locations inside the flash file that it is simpler to supply a template rather than code it all in from scratch.

Other requirements are usually around the supported formats (the restrictions for which can be gleaned from the Component package: and by looking inside the ad server interface). Sometimes the ad server specifications restrict the number of supported child files, and the file weight of the overall ad unit, so rarely will a rich media creative exceed ten child files, and file weights of over 30mb are still a rarity (although most ad server can handle more if pushed to do so). In fact the best response to a query about specifications is to look at the Publisher specifications first, as these contain much tighter restrictions than ad server specifications themselves.

Publisher Rich Media Specifications

User experience on Publisher websites is greatly impacted by the speed at which users can navigate and load the site. For this reason, Publishers spend time carefully choosing what the file size limits should be of objects and content pulled into a page (as this directly impacts the page load time). Standard flash ads and Rich Media alike need to be limited to certain available ad dimensions and file sizes and Publishers release "Publisher specifications" to media and Creative Agencies to keep the user experience at its optimum (most Publisher specifications are publicly available).

It is imperative to have the creative developer build the creative to these limitations so that ads and ad tags are not

rejected when tested by the Publishers ad operations team, before the campaign goes live. Publishers across the board have varying Publisher specifications and this means checking before a campaign to ensure the development team have received specifications from all Publishers. Invariably to save time, Creative Agencies build to the "lowest common denominator" of specs, meaning that the most restricting specifications dictate the build, so that the same ads can be distributed to multiple sites without issue. This is all the same for standard display as it is for Rich Media but Rich Media specifications will be more complex and flexible than those for standard display; particularly with regards to ads utilising methods that do not disrupt the user experience. The most common of these is the parent and child file method; known as the "polite" load.

3. Parent and Child files - the polite load

Parent and Child files

Unlike the vast majority of standard banners (which consist of a single ".swf file"), Rich media ads can be made up of many different files. The parent or master file loads first on the Publisher page (containing the core code required for the ad to pass the rich media specifications of the third party ad server). A rich media ad can be made up of a single master file and a number of child files, where the master file makes a reference to the child files by name so that their content can be pulled into the ad as it plays.

For example a master file containing the controls for a video player, references a child file called "video.flv" (where the child file just contains a video) and could not be played in the browser without the controls and code from the master file. Child files come in many different varieties; from different file types of video to PDFs, word documents, sound

files and XML files. When the Creative Agency uploads the multiple files to the ad server, the ad server uses the names of the child files to link them to the master file; commonly dictated by an intelligent method of pairing, such as by uploading the files together or into the same storage folder in the "creative library" section of the ad server. More advanced members of the Media Agency can take a look inside the parent file to determine which order the child files are referenced in (this can be done either in the Flash authoring tool or in various freeware programs designed to peer into the code directly). The order of the child files, just like the naming of them, is imperative to have the ad working correctly. Sometimes a trafficker uploading a rich media ad into the ad server will see an error when a child file is missing and opening the parent file up to see which file is missing in the order of files, will reveal the file's name and allow the trafficker to chase up the creative developer for the absent file.

Polite Download

Why have one parent file and many child files? Why not simply have a single parent file and be done with it? Looking back to the Publisher specifications and the discussion about the user experience; polite download means that the master file loads first, usually with the rest of the Publisher page and then the child file "politely" loads in afterwards (often as a background process), so that the user does not experience a slow page load time. A common polite download situation would see a master file of 40kb and a child file of around 2mb. The polite load prevents a disruption to the user experience and the effect is often undetectable by the user. An example would be instance of an in-page video ad where the parent file is simply displaying the holder image of the first frame of the video file with a play button watermark over the top. When the user clicks the button the video begins to play but the buffering has already started and

completed before the user hits play so there is no disruption to the video. Video and polite download Rich Media ads commonly utilise this functionality but the parent files of expandables and floating ads would be found to contain references in the code to the ad's rich media "state"; an essential concept in Rich Media reporting.

4. Expansions and Floating -the declaration of state

For both expandables and floating ads the more advanced functionality is enabled owing to the presence of what is referenced to in the development and build cycle as the "state" of the ad. (This is not to be confused with "State-based ads", which allows the user's interaction with an ad to be remembered for subsequent encounters with the campaign) Hidden away in the code of a Rich Media execution for both these two formats is the ad's state. When the ad changes its state, further actions can be undertaken by the code. For example with expanding ads, the ad has two key states: Expanded state and Contracted State. The creative developer can insert commands into the core code of the ad, to run as the it enters or returns to these states, such as the command to being playing a video or sound file. This freedom is also granted to the trafficker who can (as with standard display ads) piggyback extra code to execute on a change of state such as the firing of a piggybacked pixel or conversion tag (see Chapter 4). States are important for reporting as well, since the change from one state to another fires off default interaction tracking such as counting the number of times the ad has been expanded.

5. Interaction tracking: Default and Custom Events

Rich Media can support the tracking of interactions which means that apart from the usual tracking of impressions, clicks and conversions; the rich media file can send information to the ad server when the ad registers any action at all and these appear in reporting classed as "interactions".

Interactions can be timers, rollovers with the mouse, changes in the ad's Rich Media state, clicks, user gestures and the pressing of keyboard buttons. Rich Media templates supplied by the third party ad server to the Creative Agency will pre-code a number of default interactions. A good example might be that for a video ad, the ad server records when the user is 25% through the playing of the ad and then 50% and 100%. Each of these interactions will have its own column in ad server reporting to aggregate the number of times they occurred; which assists in understanding creative performance. For example with the trailer for a film, a creative that is never recording "100% plays" may contain video that is too long. Other default interactions include total number of expansions and a timer for interaction duration. It's worth noting that default interactions are not industry standardised meaning that there is no list of all default interactions that all ad server and Rich Media tech providers will supply.

It is therefore worth checking with the third party ad serving provider exactly what comes as a default (an example Rich Media report will display this). Furthermore the methodology behind default interactions is different. One to watch out for is knowing when the interaction timer begins: some ad servers do not track "accidental" rollovers as part of the ad engagement while others do. This can explain quite varied differences between interaction rates between

different technology suppliers. A useful default interaction for Advertiser utilising TV as well as online delivery mechanisms is the "brand exposure duration" metric. This can make a rich media ad comparable to an ad for TV campaign for the marketer on the Advertiser side, where the total viewing hours nationwide are captured and tallied. All such metrics are simply default interactions and appear in ad server reporting (see Chapter 5).

Thankfully Advertisers are not left just with the default interactions but have the ability to inform the creative development team to add custom interactions. Custom interactions can be added by the creative developer to any part of the ad during the ad build process. This could mean the pressing of a custom button, the user hovering over the brand logo, or even the measurement of time between interactions of different types.

Custom interactions although useful have limited application so it is important not to have the developer go overboard when implementing them. During the kick off call, as the Agency or Advertiser advises the creative developer of the interaction requirements it is worth considering what actions can be taken as a result of high or low figures collected for these interactions. For instances, tracking and reporting on the number of times a user scores a goal in an interactive Rich Media football game ad will not inform either the Advertiser or Agency about user brand engagement and may only serve to append the goal scoring functionality in future iterations of the updated creative.

6. Interactions as additional clicks

Custom interactions add almost unlimited flexibility when it comes to tracking. Another key component for Rich Media

ads is that they support multiple click-redirects. So, different elements of the ad can click through to completely different Advertiser web pages. The ad server allows all clicks taken by the user in this manner to be recorded as clicks but reporting can break out the increment of clicks by their click through URL to discern which page received the most traffic while still accurately reporting on the "total number of clicks". More on this in Chapter 5 - Reporting.

7. Building Custom Rich Media ads

It is all very well and good standardising the most popular and well used formats but the whole concept of Rich Media is based on creative impact. To stand out in a crowd one must look different and in Rich Media, such an achievement sometimes literally means 'thinking outside of the box'. To achieve executions that look incredibly impressive, such as those that shake the whole page and seem to dismantle elements that remain functional (such as search boxes and navigation), the creative developer must work alongside the Publisher. Often the ad serving provider will liaise with the Publisher directly to cast aside the Publisher specifications for a single high impact execution that, like a firework, is very dramatic but shorty lived. Most home page takeovers utilise an element of customisation which ads longer to the development time of the ad and will come out pricier than sticking to standard formats.

Once the creative developer has completed the work of adding functions to the design team's creative, the ad (just like with standard display) is ready to upload to the ad server. This can be undertaken in most instances from within the ad server supplied components inside the Flash authoring environment. Unlike standard display ads, Rich Media ads tend to be made up of several files (at least one parent file and one child file). Some ad

servers allow developers to zip all the files up together in a *.zip file or folder before uploading them to the ad server for the trafficker to work on them. Before the trafficker completes the trafficking process, the complexity of the Rich Media ad requires it be more thoroughly tested and QA'd.

QAing and testing Rich Media ads

Building rich media ads to specifications takes some time and some Creative Agencies have not worked with some ad servers before, and so will take time to install and get used to their components. The kick off call, as described earlier in this section helps ensure creative developers are fully trained before undertaking the work. However, as with using anything for the first time, progress will be slower than allowing the Creative Agency to work with tools that they are already familiar with.

This is an interesting discovery for ad serving providers because last minute campaigns require fast creative processing and Media Agencies and Advertisers will sometimes defer to the Creative Agency's preferences when selecting an ad serving technology provider. This decision can be solely based on the usability of the tools and how comfortable the developer is with support received from the ad server directly in times of trouble.

Development time varies dramatically, but a single ad concept with rich media elements from scratch never takes less time than a single working day (if you find a Creative Agency that can do this; check the quality of their work and then stick with them!) In fact the longest part of the process for any rich media campaign setup is the development/coding stage.

The Creative Agency is expected to test their own creative locally and potentially even in a simulated Publisher environment on the third party ad server (something that the Media Agency can also do) before the ad is invariably tested manually by the third party ad serving support team as well.

Once the creative passes these checks (a process called QA - for Quality Assurance), The Publisher will also test the ad via the received ad tags before putting the ad live. The ad can fail to pass QA at any one of these checkpoints either for specification failure, or functional failure and the creative is pushed back to the developer to amend or rebuild; sometimes with assistance from the support team at the third party ad server. Ads can also fail even once a campaign goes live despite rigorous checks so QA becomes ever more important.

Trafficking Rich Media ads

Apart from the extended trafficking time required for QA, rich media ads are treated the same way as standard display ads, so that once the ad is created in the ad server, it can be assigned to a placement and the tags can be sent to the Publisher.

As Rich Media ads can be more expensive than standard ads to serve through the ad server, the temptation is for some Media Agencies is to go directly to the Publisher. The danger here is that the Advertiser loses the ability to de-duplicate the user to understand their life time interaction with a brand when the wider picture of a pathway to conversion is considered. In addition independent third party reporting is unavailable and complex interaction reporting is not commonly available in publisher side rich media serving technologies.

Creative Agency and cross Publisher support would also vanish meaning that small problems can quickly escalate beyond repair. Third party ad servers conveniently offer tools to make sure that Rich Media ads are not wasted such as through on-interaction retargeting (Chapter 7) or with ad controls that ensure Rich Media is reserved, deeper in user sequencing or manipulating frequency capping to ensure Rich Media ads are shown frugally to users or

displayed at more effective times of the day. No one is denying that Rich Media is a cheap alternative to other channels but the value of the impact can be harnessed through effective trafficking by manipulating clever ad controls. As explained in the previous section, the experience of the trafficker would add enormous value to media budgeting by being aware of all of the ad server capabilities.

Publisher Troubleshooting

As mentioned even after QA and trafficking, the Rich Media ad could fail on the Publisher website owing to an oversight in the technology or human error in setup. In such instances technical support are bought in via the third party ad serving provider to work with the Publisher to fix the problem. This is an additional operational cost that Publishers must bear in mind before offering to allow Rich Media ads to be implemented on sites to utilise greater income on the inventory.

> But how after so many checks can a rich media ad fail? The most commonly reported types of failure on a Publisher site are to do with how the different elements of the page are layered, such that the ad appears underneath or on top of the web page content. Generally speaking, the cause of the problem is a wonderfully exotic CSS element called the "Z-index". Z-index refers to the layer that each part of the web page lives on and there are thousands of such layers that could potentially exist.
>
> Although the ad is designed to appear on the upmost layer, sometimes this does not occur and the ad server may have to use a few JavaScript tricks (where able), to influence the layering of the elements around the ad so as to push them down on the Z-index and raise the ad upwards. This is made available by piggybacking additional code onto the existing ad or even inserting it into the creative itself. Publisher troubleshooting, although time sensitive for live campaigns,

can be a process of trial an error; especially with limited Publisher support, so patience is mandatory while a conflict is being resolved.

Often the support team handling the request will require that the ad tag be loaded in the live page or an exact replica inside a publisher "staging environment" so that through trial and error with code, the issue can be resolved and the campaign put live again. Rich Media ads put live on " blind networks" face the troublesome conflict of having to supply a URL to the support team to resolve the error while crucially not exposing the site name to the Advertiser or Agency. For this reason it is not recommended to run Rich Media through such agreements.

Due to the open ended nature of Rich Media (to express creative flare through technology) there can sometimes be unforeseen technology conflicts which may result from updates to browsers, updates to the core Flash player technology from *Adobe* and other anomalies in JavaScript and Flash elements on the same page. All of which can take a thorough investigation to resolve. Thankfully Publisher troubleshooting issues rarely occur and the Advertiser's return, for high impact formats, are both rewarding to the user visually and facilitate campaign goals for effective reach and greater interaction for the campaign.

Rich Media for Mobile & Rich Media for Instream

Thankfully once one has mastered Rich media, most of the other capabilities of the Display channel seem like child's play. Beyond Display, Rich Media has begun to penetrate the channels of Mobile and Instream too. These nuances are covered in the remaining sections in this chapter. But before leaving Display advertising behind completely there remains an elephant in the room, one which has allowed giants among web technology companies to thrive. This would be the ability to create display

ads which are Dynamic in nature, circumnavigating the need to have the creative developer build quite so many ads. The next section focusses on what is referred to as DCO Channel or "Dynamic Creative Optimization" and it unlocks a universe of true creative optimization.

Chapter 3 - Section 4
DCO

- Standard XML Display
- *Figure 3.4.1*
- *Figure 3.4.2*
- DCO Design Phase
    - Creating the dynamic template
    - *Figure 3.4.3*
    - Dynamic Creative Planning
    - *Figure 3.4.4*
- DCO - Ad Creation Phase
    - Single Versioning or Manual Creation
    - Excel Mass Versioning
    - Mass Versioning from a Feed
    - *Figure 3.4.5*
    - Using a web service to restrict load on the Advertiser hosted product images
    - *Figure 3.4.6*
- DCO - Assignment to Placement Phase
    - Placement vs. audiences
    - *Figure 3.4.7*
    - Targeting different Versions to different Audiences
    - *Figure 3.4.8*
    - Connecting DCO with Audience Buying
- DCO - Optimization Phase
    - Auto-Optimize Vs. Manual Control
    - Figure 3.4.9
    - Auto-optimization as a "Black Box Technology"
- Multi-variant testing

Standard Display and Rich Media ads could be described as "Static" ads so that once the creative developer has completed building the ads and has uploaded them to the ad server, the content of that creative does not change. The elements or assets that make up the creative are rigid and cannot be moved without republishing the master file. If an Advertiser or trafficker wanted the ad to look different, for whatever reason, the creative developer would have to adjust the creative on their side and then upload it to the ad server so that the trafficker could replace the existing with the new ad. This process can be a costly undertaking for Advertisers with regards to creative costs and for almost ten years ads created for ad serving for display existed only in this construct.

Creative that can be visually adjusted with content controlled by the trafficker and not by the creative developer means that costs all remain on the media side (Agency and Advertiser) as opposed to on the creative side (creative developer and Creative Agency). Controlled costs or lower costs are more attractive industry wide and mean more can be spent on buying the ad space. The creative can be designed so that its elements are more fluid and not "static" and this is called "Dynamic Creative". The utilisation of dynamic creative for the purposes of improved campaign performance in the third party ad server is called "DCO" or "Dynamic Creative Optimization".

DCO began life when creative developers realised that if they used an XML file type when they created ads that it would save them a lot of time if they wanted to make changes remotely later without having to republish the flash files . This section starts off by looking at these early XML driven dynamic ads before focussing on DCO, looking at the capabilities of this 'channel' and the bright future which lies ahead for its application across Display, Mobile and Instream.

Standard XML Display

As a .swf file is near impossible to break into without heading back to the original .fla file and republishing the ad, creative developers found it too frustrating and time consuming and resolved the access problem using an XML file instead. Unlike .swf files, the XML file can easily be edited in a program like *Notepad* or *Word*, so a creative developer's flash skills and premium access to the flash authoring tool are not requisites to changing the content of an XML file which in turn is used to make adjustments to the .swf file. By coding the .swf file to look for its instructions in a neighbouring XML file, the XML file is suddenly an easy-access gateway into the ad to change the content; be it a price-point, a tag line or even a background colour. Using this method a creative developer could build one piece of creative with an empty text field; the creative could reference an file called "content.xml" which could contain the English-written text to populate the text field, written out in full.

Providing these two files to the trafficker (advert.swf and content.xml) they could upload both files to the ad server for one campaign and then replicate the two files and replace the content inside the content.xml with the same content written in French before uploading the two replicated files into the correct place for the French campaign.

Before XML files came into use the creative developer could have been paid to produce two pieces of .swf creative instead of the one .swf produced today, thus creative costs come down.

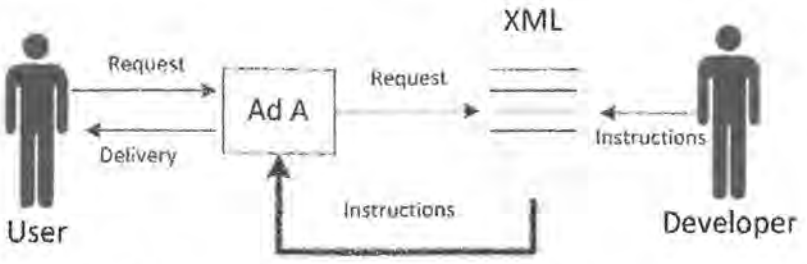


Figure 3.4.1 *the developer can update the ad indirectly with an Ad purpose built to receive instructions via an additional XML file. When the user's browser requests the ad from the server (Ad A) the ad looks for its XML file which provides the latest instructions to update the Ad. The user then gets delivered the updated Ad.*

This example suggests that an XML file can be populated with content but it can also be populated with URLs or references to the location of images or additional content. An .swf file could appear to contain a series of images but behind the scenes, the .swf file is referencing a file that might be called pictures.xml and pictures.xml could contain a list of the location of all the images set to appear the series in the ad. To change the images displayed in the ad, the traffickers could edit pictures.xml to introduce new images or remove old ones. In addition a .swf file could be coded to reference an additional section within pictures.xml that contains a number used to control the speed at which the images in the series rotate (perhaps the higher up the number, the faster the rotation). An XML could contain a whole series of remote controls to adjust just about anything going on in the ad. Each set of controls exists in separate section of the XML file called a "node", referencing the right node in the accompanying XML file maps the control to the function.

Even the XML does not need to sit alongside the creative it is updating. The XML can be completely remote so that in an ideal world an XML file can in theory sit on the desktop of a planning manager from a Media Agency and they can update the ad from

their local machine. (To perform this feat the creative needs to have the URL of the XML setup to look for the user's local machine as a unique IP on a network). More commonly the XML can sit on the Advertiser's own site but in either case there are security settings in place to ensure that the XML cannot be remotely accessed or adjusted. A "Cross-Domain" file is a very small file that sits on the Advertiser's top level domain and grants access to the third party ad server to allow it to grab the XML file to make changes to the creative. With this direct connect to the XML file, creative can be changed in real time (and at scale the change will be live everywhere within 15 minutes thanks to ad server ad replication).

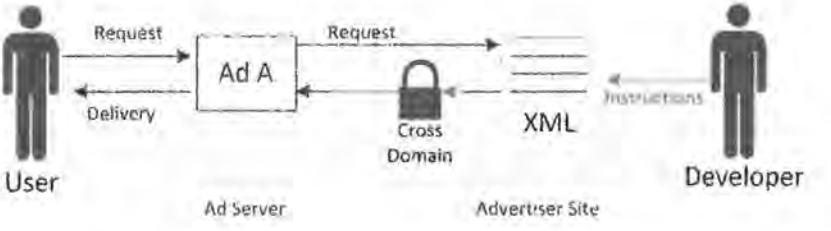


Figure 3.4.2 *With a cross domain file in place, an ad can securely call an XML on an external server (such as on an Advertiser site) and as long as the ad server has been granted access, the XML is free to deliver the information back to the Ad without opening up an unsecure communication channel to the external server.*

As great as this sounds, creative developers very quickly encountered a problem; changes made to the XML file could not be previewed in the flash ad before the amendment was permanently added. Without a preview, new text could bleed into the rest of the creative unintentionally or images could be referenced that don't exist and incorrect settings, set remotely, could cause the ad to become unstable. Furthermore allowing

the XML to be changed so easily meant that it was near impossible to track changes made to the ad (when republishing from a flash file, developers typically append the name of the file, creating a new version of the same ad). Whole tools were required for developers to really make the most of the XML remote capabilities and these were to arrive on the scene in an interface as DCO.

DCO - Design Phase

Creating the dynamic template

The modernised workflow for Dynamic creative requires that the creative developer only produces one creative for each ad which can be better off considered a 'template' creative (containing all of the ad possible designs and functions). A separate template would be constructed for each of the required placement dimensions for the campaign. The template essentially a blank canvas which contains elements or assets which are dynamic in nature, meaning that they have the capability of changing the way they look and behave.

To build such creative, the creative developer utilises the third party's component pack inside the flash authoring environment to create dynamic text box's and empty image slots in which text of a fixed length or images of a fixed size can be dropped into at a later date. The dynamic elements can technically be anything that can be built in the Flash environment which includes text, images, videos, sounds, special variables, files, strings and integers. The rich ability of Flash to offer the flexibility to work with so many multimedia formats means Creative Agencies still have the opportunity to get very creative despite the fact that they may only be producing one template.

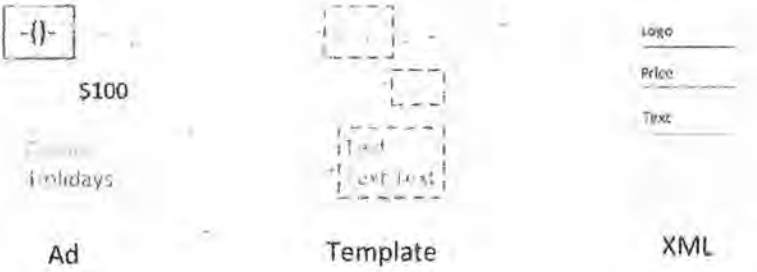


Figure 3.4.3 *An Ad can be re-designed as a template by making the areas that were previous static assets into dynamic areas which can be populated with similar content. The template, when plugged into the XML describing the location of static assets or static content can be used in conjunction with the template to reproduce a new Ad. This is called an Ad Version rather than an Ad since it is a version of the original.*

Dynamic Creative Planning

At this early stage planning is very important; for instance, how many dynamic elements should the creative developer put into the template? If one element is reserved for a tagline as text based copy, one element for the price point and one for the image there are three elements which can be changed utilising the XML, but is that enough? Some creative developers go the extra mile by making dynamic the image placeholder for the Advertiser's logo and even the small print terms and conditions.

Planning ahead about which elements will be created dynamic is essential, since once the template is created, it could be dynamic enough to be used for an almost infinite amount of time, even reused for several campaigns, without ever reaching creative fatigue by the browsing user.

Planning and identifying the dynamic placeholder is the first step, the second is to consider the possibility of text copy leaking into the rest of the creative. If copy is replaced with a longer piece of copy later on, or if images need to be used that do not fit the dimensions of the ad, then the creative planning process must make such considerations and the dynamic capabilities of DCO creation will accommodate such needs. For instance, the developer could add code to the template so that the creative automatically resizes text or images to fit the dimensions of the template.

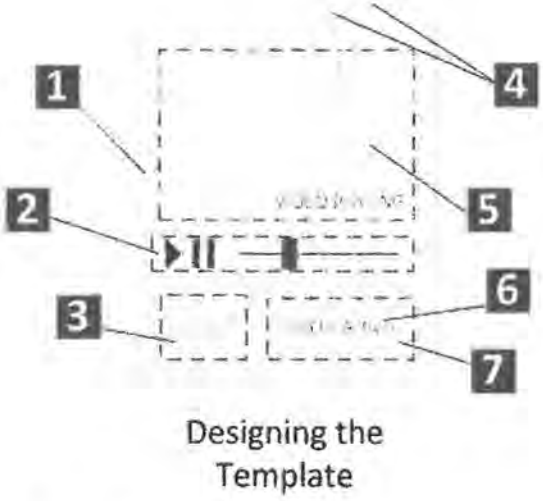


Figure 3.4.4 *Dynamic Templates should be designed for the long term so that an XML can be mapped to any of the Dynamic features and thus control them. Should the Advertiser wish to control anything that was not made specifically Dynamic later on, a new template needs to be published by the creative developer so building for most eventualities makes sense. Each of these dynamic elements are: (1) The background colour of the version (2) The player controls; perhaps there are several sets of controls to choose from loaded up as an additional dynamic SWF file (3) A*

dynamic spot for the logo for those particularly long term templates that might outlive the brand logo. (4) Particularly innovative developers that might animate part of the ad will consider adding additional tweens or paths that can be activated by the XML later on (5) A dynamic video (6) The call to action itself could be adjusted for the actual click-through URL or the text (7) A dynamic field for the font type so that it can dynamically adjusted via the XML later on.

Once the template is finalised the creative developer can upload it like any other ad to the third party ad server. The ad server will detect the use of the dynamic component used by the developer to create the dynamic elements and the trafficker working in the third party ad server interface will be given the option to decide how many "versions" of the creative to create. Third party ad servers and specific dynamic creative technologies can support the creation and management of hundreds and thousands of versions: true scaling for standard display. Subsequently these versions can be assigned to placements directly (just as standard ads would normally be assigned) or for scaled personalisation and targeting versions can be assigned to specific audience types.

DCO - Ad Creation Phase

Versioning

Versions of ads adds two new levels to the hierarchy in the third party ad server. Under the ad level exists the version level and beneath this, the asset level. (At time of writing, ad servers could not offer reporting or tracking capabilities below the version level to the asset level but the industry is expected to adapt quickly to accommodate the asset level as well). To focus now on the version level, there are three common ways of creating a version (also called "variant") ad from the template creative supplied by the creative developer:

1. **Single Versioning or Manual Creation**

Some ad servers which support the creation of Dynamic Creative, provide an interface on the platform which allows for the basic dynamic template supplied by the creative developer to be adapted into Dynamic versions, version by version. Such interfaces are essentially creative tools which can be used a trafficker to allow for the input of copy into the new ad version, the ability to upload or reference an image from within the version, insert a price point and also to change the colour of the text, the ad background colour, the size of the text (in some instances to even spell-check the copy), before the version can be previewed and saved.

The preview option is popular with Advertisers who like to proof read or view every version of each ad that potential customers will see, so that they can provide their approval before it goes live on the web. As with the preview option for standard display ads in trafficking standard display campaigns, the approval option was removed from many ad servers as the Advertiser approval chain became a bottleneck to campaign delivery timelines.

Despite a variety of supporting features, the essential point of manual version creation or "single versioning" is to undertake the design and build of a manageable number of versions, typically for a very small campaign or at least a campaign with very few ad versions. Single Versioning is not a scalable method and has its merits when an Advertiser is playing with DCO for the first time. More seasoned Advertisers and Agencies move things up a notch to "Mass versioning".

2. ***Excel* Mass Versioning**

Ad serving providers have listened to big Advertisers and Agencies running large scale campaigns with many placements and sites by providing bulk tools which utilise the easy replication of rows of data to scale. The same methodologies have been extended to Mass Versioning in the absence of a better designed UI for creating a large number of unique ad versions. Using a spread-sheet layout inside the ad server interface the trafficker can begin to make new ad versions, row by row from a single supplied creative template. The columns are already pre-defined by the ad server interface which identifies the dynamic elements of the creative and converts them into columns. Each new row represents a new version of the ad and each cell can therefore be customised with a value (such as a whole number for a price point, a text for the copy and URLs to point at external content like images). A spread sheet view quickly allows for the update of many existing versions and adds 'copy and paste' capabilities to create new versions. As with Single versioning there are checks in place for spelling as well as preview functions for each of the rows so that on mass, preview links can be generated and sent over to the advertiser for them to approve.

3. **Mass Versioning from a Feed**

The creation of hundreds of ads is a common demand for Advertisers and their Agencies where a vast product line is being supported. To build separate ads features different products in each one is completely unattainable for the likes of retail Advertisers which is where DCO can come to the rescue. Such product catalogues are usually already aggregated on the Advertiser website and therefore a product list or "feed" exists (behind the scenes of the Advertiser website). Such a list might even be as complex as a database with information that might be updated very

frequently such as the colour of the product in question and the number of items left in stock.

In turn this feed allows the CMS to dynamically generate web pages and the same method can be used to build out versions of ads in their thousands. The ad server contains controls in the interface which allows the trafficker to connect to an external database or feed and generate ads from it. The controls also permit the trafficker to map variables such as the stock control to the version and create a rule that removes the version from ad serving when the product runs out of stock. If the price point or stock number is manually changed in the feed it should reflect in the served version. In general the available DCO solutions coupled with ad server in the market range in complexity. *Flashtalking* for example allows the trafficker to browse to the Advertiser site from within the ad server interface to select a product feed that might already be displayed on the Advertiser site. *DoubleClick* allows the trafficker a much more integrated setup when connecting to Advertiser feeds that have been setup using *Google's* own stock control database tools.

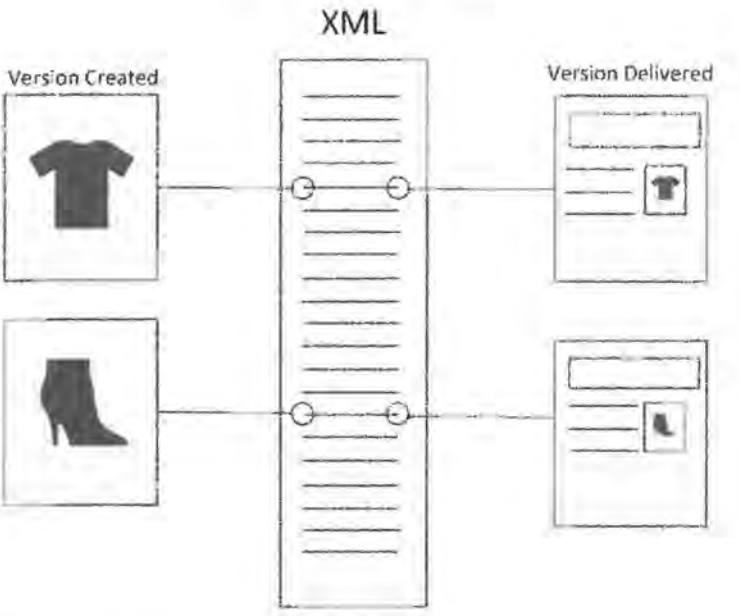


Figure 3.4.5 *a feed (such as an Advertiser's product feed) can be used to generate a mass number of versions. Versions delivered to Publisher sites are those generated from the feed*

Sometimes product feeds and databases are complex or partitioned out and require ad serving provider support to build out a new feed from which to create versions. In addition if product feeds are likely to change frequently, the DCO tool needs to be configured to 'check' the feed with the same frequency so that dynamic variables like price changes are reflected as close to real-time as possible.

Using a Web Service to restrict load on Advertiser hosted Product images

A mass versioning campaign with a product feed will typically pull the images of the product into the ad version from the Advertiser website directly and not necessarily the third party ad server. This is done by having the URL location of the

image included in the product feed and mapped to a corresponding dynamic image placeholder inside the master ad template.

Hosting images and content is a common query that comes up in the DCO setup. After all, with so many images to reference thousands of products; can the Advertiser not simply direct the ad version to retrieve the image direct from the Advertiser site each time the ad is called?

A couple of considerations need to be taken ahead of time. The first consideration is that an ad server is a piece of technology that is purpose built for mass traffic. If the versions and images themselves are expecting to be referenced by every end user on the same scale then the Advertiser site needs to be setup to handle just as many ad calls or risk being taken down by users, demanding content from the Advertiser's own servers.

To get around this, the images themselves can be pulled from an Advertiser site and cached by what is referred to as a "web service" technology sitting between the Advertiser site and the ad server. Images are cached by the web service which takes the risk of high volume traffic away from the original host URL. Caching images in this way comes with its own risks so it worth fully exploring this topic with the ad server in question before setting up and running such a campaign. In such a scenario, an adapted feed would need to be appended to the existing product feed so that new URLs for the product images are referenced by the dynamic field in the ad template.

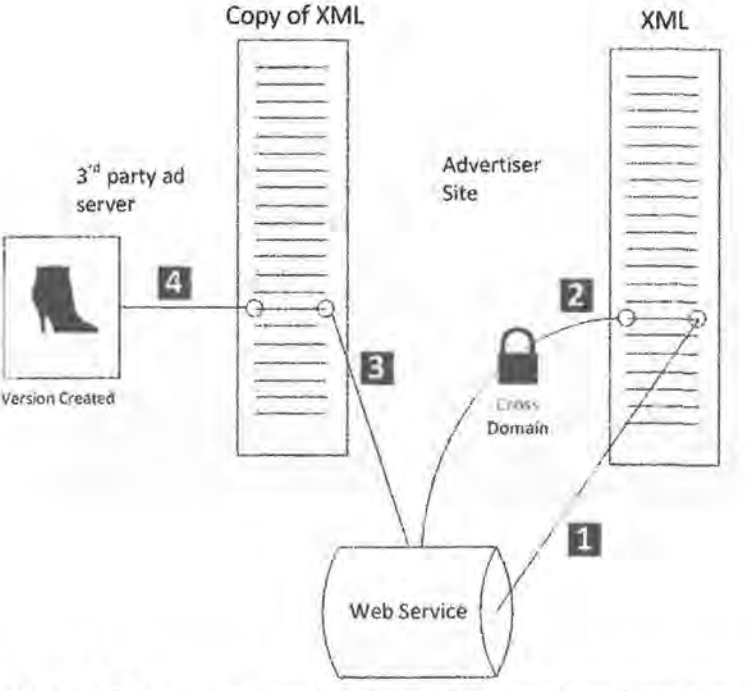


Figure 3.4.6 *Hitting product feeds frequently to keep the versions up to date can be a heavy strain on some Advertiser sites or servers. To get around this problem a Web Service can be setup to replicate the product feed on the ad server. The web service would make a request the core XML on a limited frequency (1) for example once per hour instead of for every ad call. The request returns a response from the feed via the cross domain (2) for security purposes. The web service then updates the various part of a copy of the XML (3). When ad versions are requested by the user the versions are up to date which what is in the copied XML on the ad server (4).*

Feeds themselves are also rarely ready to plug directly into a DCO setup on the ad server side so Web services can be used to scrape the original product feed at different times and act as an adapted feed for the creation of the DCO versions. Feed Scraping is not usually a free service so it is worth an Advertiser or Agency fully investigating feed adaptation and scraping when working out how much it will cost to create and setup a mass versioning campaign.

As with legacy XML creative setup, if a feed sits inside an Advertiser's domain or firewall it will require the presence of a cross-domain which accepts the calls to and from the ad server to work correctly (adhering to the required security protocols to accept the ad server traffic). The use of a web services eradicates the real-time element that was so attractive about XML dependent creative, the amendments to the version will only change when the web service updates a copy of the feed or re-caches the stored images. This process of scraping at the time this book was written was around one hour maximum, so a change to the feed, should It change a version (such as with an update to the amount of available stock on a product) would only be reflected up to an hour later. Some ad servers have even longer time lags on this.

With the freedom to create so many versions, comes the responsibility of identifying each of those versions individually. Each version is assigned an ID in the ad server for reference purposes, since the only other real unique identifying marker is literally the unique combination of dynamic variables that make the version unique in the first place. If a version is changed or adjusted, its unique combinations of dynamic variables no longer matches the original configuration and this is a problem for reporting. In truth any adjustment of a version should create a wholly new version so that the versions can be discerned in ad server reporting. In turn the creation of a new version as with the creation of a new ad requires a certain amount of re-assignment to placements or audiences which on a large scale (such as using a version across a large number of placements and sites) can prove to be a trafficking headache. The most recent upgrades to DCO interfaces for third party ad servers are building in the functionality and logic to cope with assignment on this scale such as by associating versions that have been adapted from one another, together.

DCO - Assignment to Placement Phase

Just as with standard ads, ad versions can be assigned directly to placements. Once the versions have been created in the ad serving interface the trafficker ensures there is the correct configuration of placements, sites and delivery groups in order to attach the ad versions.

Placement vs. Audiences

Being able to scale creative means that very deep personalisation becomes possible, since a version of an ad can be created which has more relevance to a user, than a run-of-the-mill standard display ad (designed for mass reach) would have. Relevance generally brings an improvement in campaign performance especially when DCO is combined with a source of audience data, which brings about highly customised targeting (see Chapter 4 - Audiences and Targeting).

As explained, inside the ad server the trafficker has the opportunity to assign ad versions to placements (as is possible with Standard display and Rich Media) but with DCO, the capability for assignment extends so that across placements the same version can be displayed to multiple audiences.

By assigning versions to audiences the trafficker exposes a single ad version potentially to a much larger group of users that might be exposed to a single placement. This is particularly true when a media plan features a large number of sites and within that a vast number of placements.

All users arriving at a single Publisher site and seeing a single placement will be a smaller number of users than an audience of users encountering ads across several placements on several sites. Since an ad server can identify users in an audience group, whenever the ad server serves to the audience member, their

presence in the audience can over-ride the default capability of serving the ad designated only for that site and placement.

This will only work if the definition for an audience is fairly broad. If the audience is too granular (such as users that live in a certain Zip Code who have visited an obscure page on the Advertiser site) then the rules setup by the trafficker on the placement will not find the matching audience and will therefore serve a non-targeted ad or ad version.

Setting up thousands of very specific audience groups with corresponding ad versions is very time consuming and if only a couple of users ever encounter a version, there will not be enough data to suggest that the versioning of the ad is having an influence on their decision to click for instance. If a couple of users also click on a version which is almost identical, how would one workout which creative version is the best performer?

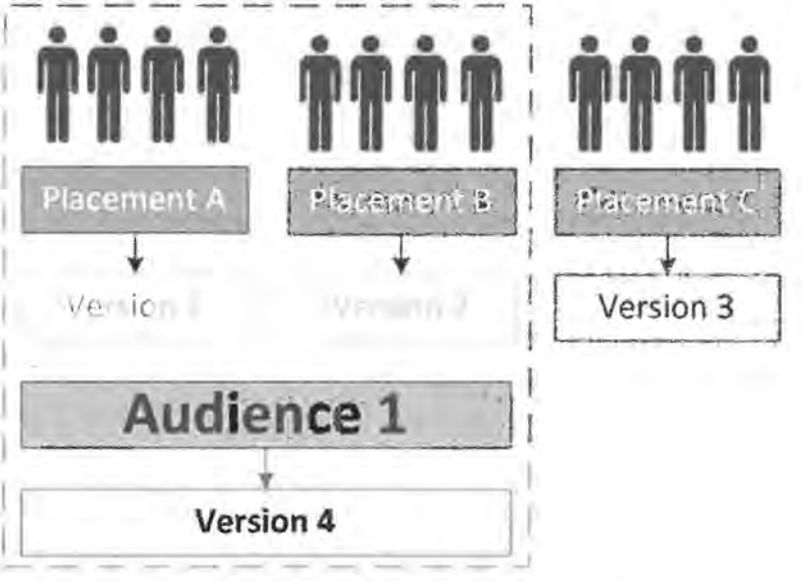


Figure 3.4.7 *if versions are assigned to placements (where one placement is on each Publisher site for instance) then the number of individuals seeing each version will be low and deciding what*

the best version is becomes more challenging as the sample is smaller. Apply a version to an entire audience and a version is shown to a wider group. This is not always the case (sometimes audiences are smaller than the traffic going to a single placement) but to get the most meaningful data set, the version should be applied so that it is seen by as many users as possible with as few versions created as possible.

Targeting different Versions to different Audiences

As Mass versioning allows for a vast combination of dynamic elements to be turned into ads of their own, the next step is to involve the media planning team from the Agency to decide which of those ad versions are suitable for which audience type (since the assignment to placements alone is circumvented). This means defining the effective audiences in the ad server and then assigning the correct versions to the correct audience. The media planning team are important in this step since the media is often bought to match the audience targeting requirements provided by the marketing team.

Audiences can be assembled or defined in a variety of ways for both dynamic creative and for standard creative. Chapter 4 deep dives further into the available audience groups for targeting inside a third party ad server (something that differs from the audience targeting capabilities offered through a Publisher ad server). Of these types, the "retargeting" audience is the most common for which a single ad version might be highly relevant to a user.

A retargeting tag (as will be further explained in Chapter 7) can, for example be placed on an Advertiser's product page: invisible to the user and allow the user to be targeted again later if they leave the Advertiser site without purchasing. The result is that an ad version (created using mass versioning) matching the last product the user saw, can appear to effectively chase them around Publisher sites until they make the purchase. To enable

the chase the media planners procure the inventory required across Publisher sites. Pure DCO solutions like *Criteo* do all this legwork for the Advertiser in question, buying the inventory on their behalf with the intention to chase down such a sale, using mass versioning to keep the creative relevant.

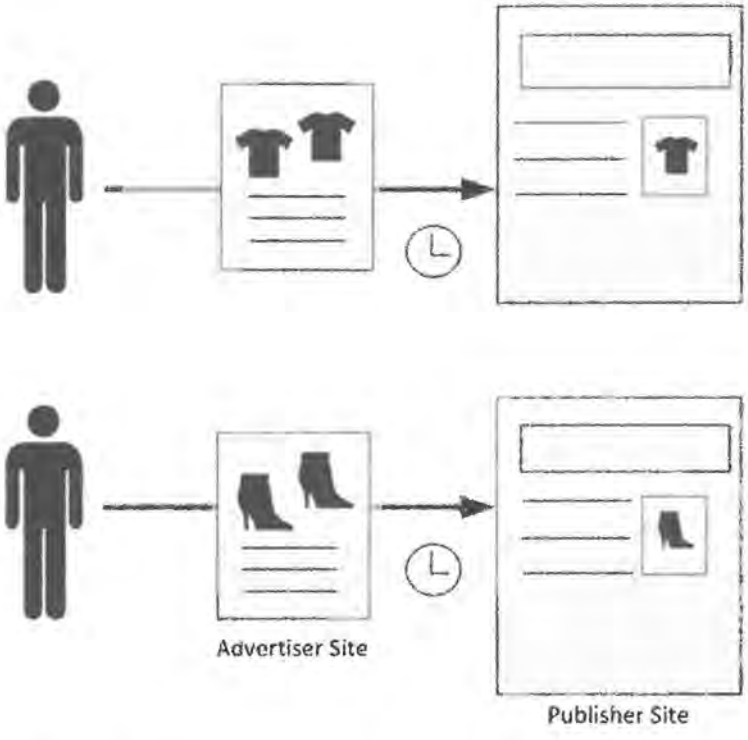


Figure 3.4.8 *Audiences visit the Advertiser site, in this example the product page where a retargeting pixel adds their cookie to a cookie retargeting pool. When the audience later encounters an ad spot bought by the Advertiser on a Publisher site, the ad checks to see if the user belongs to a retargeting pool and if so shows them the relevant version. This is all possible with standard ads but DCO allows for this at scale.*

Connecting DCO with Audience Buying

The ability to deliver a customised version of an ad based purely on a user's actions means that DCO can be used in conjunction with machine-based audience buying systems in the RTB space. This is a significantly more advanced concept, which is covered off in more detail in Chapter 8.

Pinpointing a user and showing a relevant piece of creative in the form of an ad version has spun off into two core technology features Dynamic Retargeting and Carousel Ads. To digest both of these two technologies requires a deeper understanding about retargeting tags (and though these are features of DCO ads) the detail about their setup and implementation are detailed in Chapter 7 - Retargeting.

DCO - Optimization Phase

Auto-Optimize vs. Manual control

Once the versions are created and assigned to placements and audiences there comes the question of how to decide which versions to show the end user. Although personalisation comes into the game for versions created using corresponding product feeds; in the case of building out hundreds of versions of ads that have different creative nuances (such as different font and background colours) it is less clear what version to show to who and when. Sophisticated rules are required because optimizing many ad versions is a bigger task than optimizing standard ads. With versions en-mass it makes more sense to allow machine-learning to "optimize" the versions so that those that are performing well are exposed to a wider effective audience and those that are not doing so well simply disappear. Having said this it is worth checking that the third party ad server that supports dynamic creative allow for both manual and automatic optimizations since these features are not universally available.

The control of weighting or sequencing ads, or running different ad versions at different times with various frequency caps, are options that remain open in DCO as they are in standard display and are effective tools in refining a campaign (with a small enough amount of versions for the target audiences seeing those versions to not become too granular).

Once a DCO campaign is setup manually by the trafficker; the automated approach can prove to be a powerful tool if well managed at arms-length. Auto optimization has some core facets.

1. First of all the setup needs to consider how to optimize
    a. Will the machine learning adapt according to an A/B style test (see Chapter 4)?
    b. Or use multi-variant testing (see further on in this section)

2. Then a consideration needs to be made on the core success metric which may for example:
    a. prefer a lean towards clicks over conversions
    b. or favour a higher basket value in conversions
    c. Or converters from a particular geography.

3. The last consideration is the effective pool for the optimization, so this would be to display the version or collection of versions to be shown just to users arriving at a given placement or opting for a wider group by taking learning's across a complete audience.

The trafficker selects the placement from the site in question in the ad server interface and attaches the versions in a group, after selecting the preferred success criteria from the above. When the campaign goes live, those more successful ads will be shown more frequently or receive more valuable user engagements and this can all be displayed in reporting. The machine does need time however to learn based on the user activity, which creative is the

best to display. This can take a couple of weeks or more depending on the complexity of the campaign.

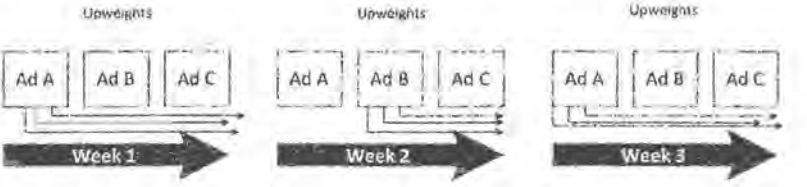


Figure 3.4.9 *Auto-optimization at work. The success metrics is the total number of clicks. High number of clicks in this example tells the optimizer what % of the audience should see the specific ad version in the following week. In the first week the weights are balanced but by the end of the week Ad A is getting the most clicks so in Week 2 most users see Ad A. However by the end of week 2, the small portion of users seeing Ad B are clicking on it at a higher rate than the users seeing the other Ads. Therefore in week 3 the up-weight is set to Ad B. By the end of week 3 Ad A seems to be getting the higher rate. The learning from this auto-optimization that Ad C can probably be dropped from the campaign as a version.*

Auto-optimization as a "Black Box Technology"

Some Advertisers ask to see the output of the auto-optimization engine making its decisions but not all ad servers offer this as an output. This is a controversial point since the learning's made by the auto-optimizer are not something that the Advertiser or Agency can retain and apply to other brands (in some instances the learning's cannot even be transported to new campaigns since the creative versions used may be different). There is merit in being able to export this data although as yet DCO ad serving technology has been slow to visualise the learning's and give the trafficker the control to over-ride the auto-optimization process to improve it based on manual optimization knowledge once it has begun optimizing.

Multi-variant testing

Get really granular with DCO and the focus is not on the creative as a complete entity but rather the tiny fragments inside the creative otherwise known as Assets. Multi-Variant testing allows the Advertiser to make decisions about the effectiveness of individual assets (for example; if a blue background on any version is always more effective than a red background). This granularity allows the Creative Agency to build a much more effective ad from the ground up without necessarily opting for DCO but requires continuous testing across a large control group to remain relevant. Multi-variant testing is popular where the creative development team is much closer to the trafficking and reporting process. The effort required to maintain something so subjective in nature is too high a price to pay for a very small improvement in conversion uplift. A direct response Advertiser may claim that every additional sale matters but if the work required to bring that sale hikes up the operational costs of developing the DCO template or the work of the trafficker past that of the gain in a couple of sales then it is not worth the effort.

We have dived into the setup of DCO campaigns that both use a few versions and those that operate across mass versions. Trafficking and optimizing these campaigns is also fairly simple providing plenty of testing and practice goes on before the campaign goes live with DCO activity for the first time and the campaign goals and audience targeting requirements are clear. DCO adapts well to both brand Advertisers (to improve and scale creative copy and content) and for direct response Advertisers (for personalised audience targeting, retargeting, and creative optimization for the purpose of driving more valuable user engagements). In the next section we begin to drift a little further a-field still from Standard Display ad serving and even beyond Rich Media into the campaign management for Video otherwise known as Instream and Pre-Roll; a certain future for an internet infrastructure that provides users with an ever more immersive advertising experience.

Chapter 3 - Section 5
Instream Video

- What is Instream?
- *Figure 3.5.1*
- Pre Roll Mid Roll Post Roll
- *Figure 3.5.2*
- Site Served Video
- *Figure 3.5.3*
- Where do Video Ad Networks add value?
- VAST - the video tag standard
    - Companion Ads
    - *Figure 3.5.4*
    - Ad Skipping and TrueView
- VPAID - Standardised Interactive Video
    - *Figure 3.5.5*
    - VPAID Authoring
    - *Figure 3.5.6*
    - DCO for Instream
    - *Figure 3.5.7*
- Beyond VAST and VPAID
- Publisher Certification
    - *Figure 3.5.8*
- Trafficking Instream ads

Third Party ad serving providers consider Instream video to be the fastest growing channel in digital. Instream Video is perhaps the distant cousin of Standard Display and a closer relative to the Rich Media channel. The trafficker actually has a much easier job trafficking Instream Video ads than the job required to undertake DCO or Search for example, the more complex legwork sits with

the Publisher.

This section will start out by explaining what Instream is before explaining how it is still in its infancy before providing an understanding about the tag structure, the capabilities of the ads in terms of tracking and reporting as well as the complete picture for campaign management from the perspective of the trafficker.

What is Instream?

Instream ads can be thought of as the evolution of TV advertising in many instances, Instream ads are literally an Advertiser's TV ad embedded into the video player of a Publisher website designed to play out video content. The ads are usually shown before the featured video content shown by the Publisher.

> Referring to Instream as VOD (Video on Demand) is a little inaccurate since VOD encapsulates the ability to stream video whenever required and would include playing video content from set-top boxes and paid for streaming services such as *Amazon's Lovefilm* or *Netflix*. None-the-less it is still common to hear those working in online advertising to refer to the Instream channel as "VOD".



Figure 3.5.1 *an example of an Instream Video ad. The Instream ad plays as a TV ad would play, before the featured video content in the same player. As with most Instream ads, some of the main video controls are disabled to prevent the user forwarding through the ad to the content faster than it is playing.*

Publishers offering Instream content are on the increase but the globally available inventory is still somewhat limited. The biggest sites for Video at time of writing include Google's *YouTube* and in the US sites like *Hulu*. All of the major broadcasters also offer in-browser video players and grant ad spots to Advertisers often in conjunction with deals designed to buy TV ad space.

To create the Instream ads, Advertisers will typically use the very same content which was filmed for TV but some are convincing their production teams to create video specifically designed for the web. The ads which are produced today tend to be fifteen seconds, thirty seconds or a minute long in play time and are converted by specialised production houses (often the function of more advanced Creative Agencies) into digital video suitable to be used online. Pure TV video is usually created at a very high definition which needs to be converted before it can be more easily transported around. There are several companies operating in the TV asset management, distribution and delivery space; big players include *Extreme Reach* and *Adstream*.

Once the video is digitised it needs to be fully adapted to Publisher specifications as one would expect from ads in both the display and rich media channels. Then it can be uploaded to the ad server and trafficked out before appearing in front of a user as a pre-roll, mid-roll or post-roll.

Pre Roll Mid Roll Post Roll

Despite the fact that video ads served before the featured video are the most commonly used Instream format (otherwise known

as the Pre-Roll), there are also in use Mid Rolls (middle of the featured video) and Post Rolls (shown after the featured video). These formats all exist in the "stream" of the video itself and are referred to as "Linear" Instream ads meaning that the ad and the video content all seem to play in one continuous stream from beginning to end.

"Non-Linear" Instream ads allow the user to perform additional actions so as to potentially extend the possible engagement period within the creative beyond the fifteen seconds of the core video. Providing that the Publisher specifications permit non-linear, this can be achieved by introducing interactive elements into the ad.



Figure 3.5.2 *Linear Vs. non-linear ads. Linear ads play through from beginning to end, only the video itself when clicked will take the user through to a landing page in a new window. With non-linear there can be additional clickable and interactive regions which, when rolled over or clicked might pause the video opening panels to a mini-website for the Advertiser inside the video player itself. The area labelled as 'F' might click through to the Advertiser's Facebook fan page for instance.*

To guarantee that videos are suitable for the Publisher or Broadcaster, media planners at Agencies and some Advertiser marketers have been circumventing the use of third party ad

serving and passing the video assets directly to the Publisher site for implementation, distribution and serving. This behaviour gives little regard for the damage it is doing to the understanding the true value of the inventory in question and identifying the ad campaign value beyond single sites or the Instream channel in isolation.

Site Served Video

Video heavy Publishers and digital broadcasters have been notably resistant in granting permission to have video ads served by an independent third party. This thinking and practice is reminiscent of the internet in the early years when many Publishers. Some broadcasters may not even be aware of the value that third party ad serving offers to the Advertiser's funding their valuable content and with a well a more well established industry in TV ad trading than in online it is no surprise to find some resistance to change.

To remind the broadcasters, third party serving allows for the de-duplication of cross-Publisher reach, frequency and the de-duplication of conversions (see Chapter 4), independent tracking and auditing as well as standardised IAB approved reporting methodology across Publishers and video players. The Publisher's excuse to site serve says more about what that Publisher is thinking and practicing than they might realise.

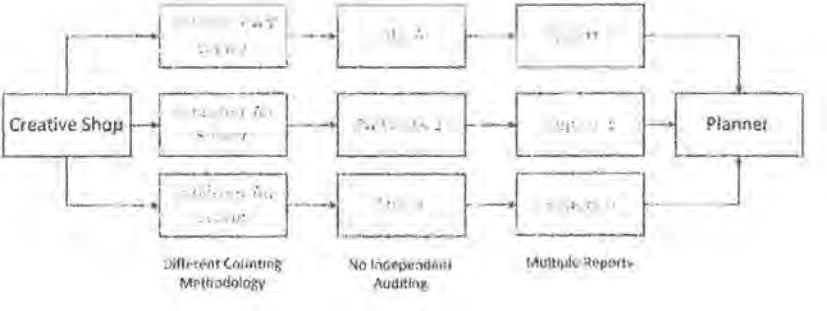


Figure 3.5.3 *Serving direct to the Publisher or broadcaster's video player without a third party creates blindness in the data and an operational nightmare to collate the generated data at the end of the campaign. The flaws are listed following this diagram: 1) there are different counting methodologies for the video because different Publisher ad servers have different counting capabilities, 2) there is no independent auditing of the collected data so it could be tampered with by the Publisher in their favour without the Advertiser knowing and 3) the process generates multiple reports in various formats that make life difficult for the planner/marketer/analyst reviewing the resulting reports.*

The trend among Media Agencies for some time was to serve Instream from third party ad servers but cut costs while still retaining some degree of transparency; by implementing tracking pixels and click trackers where the video itself is site served. This method still creates a degree of blindness and does not preserve the independent verification of the ad delivery in the ad server. A tracking pixel for instance cannot confirm if a video ad was ever seen, ever played or ever skipped. The demand from the industry has been for better standardisation and the *IAB* arrived in support of two technology standards for Instream which ensures:

- Consistent tag implementation among diverse Publishers and players
- The ability to standardise Instream metrics and core functionality among compliant video ad servers

Where do Video Ad Networks add value?

If an Advertiser chooses to run a campaign across a single network then there is no query of de-duplication because users are confined to seeing ads served from one place and tracked by a single system. For this reason Networks try to have as much reach as possible and in the Video world this is entirely possible. *Videology* is an example of a high reach Video Network. Third

party ad-serving continues to be useful on just a single network to independently verify the delivery and tracking of these ads.
Until very the very recent standardisation of Instream ad tags and technologies, third party ad serving ads was near impossible because of the many differences between Publisher players and the way they handle, process, play and track video. The standard passed by the IAB was VAST and if a Publisher video player is not certified to run VAST or VPAID ads it is best to avoid planning and buying their video inventory because it cannot be effectively managed with a third party system.

VAST - The video tag standard

VAST stands for Video Ad Serving Template, third party ad servers will output a VAST tag (which looks like any other ad tag) telling the player it is implemented in, what to display. The broadcaster's digital player will send an ad request to the third party ad server which will respond to the player by supplying an XML file that it can read. Reporting is then aggregated in the ad server as usual. VAST supports Linear ads, Non-Linear ads and "companion ads".

Companion Ads

Instream ads in the player can be accompanied by ads on the same page which tend to appear near the video content. Like Rich Media Sync Ads, companion ads are used to complement the video creative or to make it appear to be more interactive. The Publisher can use different parts of the ad tag code for the trafficked Insteam tag to assemble the companion ad next to the video content when it plays.

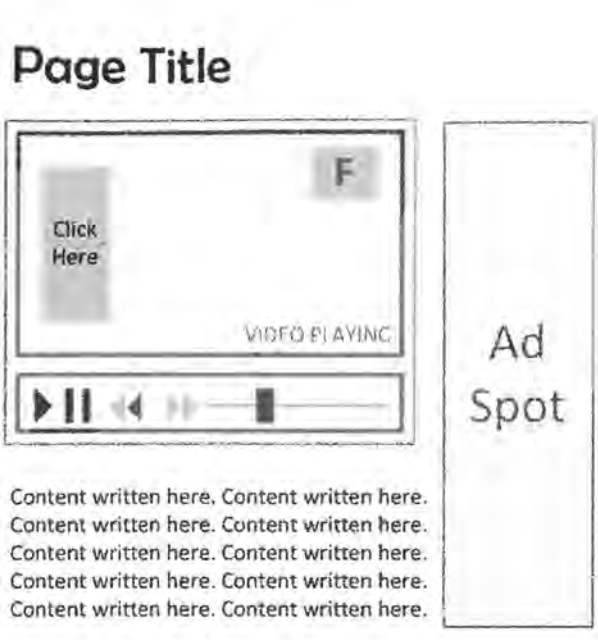


Figure 3.5.4 *A companion ad is served on to the same page as the Instream video, commonly in a reserved ad spot next to the player.*

Ad Skipping and *TrueView*

Some of the ads produced can be skipped after a few seconds while other ads play all the way through, disabling the video player controls that would usually allow the users to fast forward. In recent years *Youtube* has presented users with a button allowing the user to skip the whole ad. This technology is rolled up into a product offering called *Trueview* so that Advertisers only pay upon an entire viewing of the ad. The innovation of *Google's* skip functionality has been included in the standards for VAST 3.0 which means that increasingly, deeper interaction, starting with a skip will be available for Advertisers to implement on all ads, even outside of *Youtube*.

In addition to VAST tag standard there is an advanced player communication standard called VPAID.

VPAID - Standardised Interactive Video

VPAID stands for Video Player Ad Interface Definition and is a standardised method of communication between the video ad and the Publisher's player. Once the tag is embedded on the Publisher side, the video player calls the ad server to request the delivery of the video ad itself which is then streamed in. The Video player than performs an on-going communication between the ad and the player until the ad plays all the way through. If left untouched by any user engagement, the VPAID ad will typically play through as though it were any other linear ad but engage with it and the video might slow or pause to present new information to the user. Through engagements and video plays, the ad returns metric information to the ad server. Unlike VAST which contains around fifteen standard metrics, VPAID (like Rich Media) adds almost limitless customisation as custom metrics on top of additional standard metrics as well.

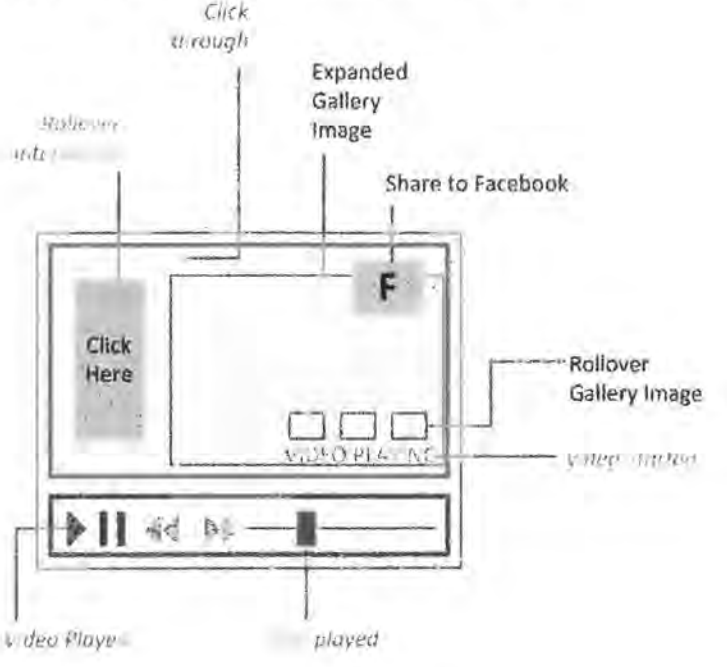


Figure 3.5.5 *A example of a VPAID ad. Collected metrics are stated in grey italics, additional functionality beyond standard linear VAST shown in black text. With VPAID there is a lot of data that can be gleaned about a user's interaction with the video ad as it plays, many of such actions can disrupt the linear play to pause the video. Functions such as image galleries might open new images in the video but on top of the playing video. Share to Facebook may open a new mini window within the player to allow the user to push a post to Facebook. Tracking for VPAID opens up the sort of options previously only available in Rich Media such as Interaction tracking and more granular video completion metrics.*

VPAID Authoring

The development of VPAID ads sounds complicated to some Media Agencies and Advertisers who are unaware of precisely how standardised these formats have become. Elements among VPAID ads are so common (such as adding a call-to-action, adding a like to a *Facebook* Fan page, *Youtube* page or *Twitter* feed) that a creative developer does not even need to be involved in the process of converting a TV ad into a fully functioning interactive creative. This is down to a whole heap of available tools which are in fact small authoring environments (mimics Flash-like authoring capabilities in some instances). Ad servers are now beginning to market such tools to Agencies and Advertisers to cut down on creative costs but the same sort of tools are extended to specialist video Publisher networks as well, these include *Videology* and *Innovid*.

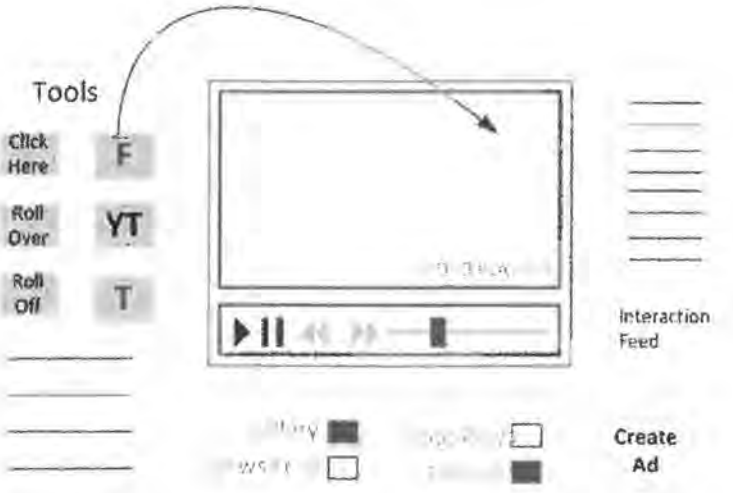


Figure 3.5.6 *A VPAID Authoring tool; allows the trafficker to transform an existing video creative into an interactive one. No coding experience is required since the interface allows elements to be dragged onto the video stage where they can be customised and the resulting data outputs can be tracked via an interaction feed panel. The completed project can be made into an ad at the click of a button.*

DCO for Instream

As with other forms of Display advertising (Standard Display, DCO and Mobile) the ad server can be used to make a decision to serve a particular creative to a particular audience type (see Chapter 4). If a flurry of audiences can be deduced, how does one go about amending a fixed piece of video to create differing versions of the creative?

The easiest solution is to use VPAID authoring tools to produce static frames in the video that appear to carry a different creative message such as different price point or written reference to a location. Some tools have emerged however, that allow the

trafficker to append the video creative by layering text, images and additional video footage both on a layer which appears to sit above the rest of the video or deeply nested in the video itself. The latter of these options produces an astounding effect, most easily implemented when layering text or images onto a blank canvas such as a wall that might appear in the background of the video.

The addition to the video appears to display in multiple frames and in multiple angles because the text or image is carefully mapped to other moving objects.

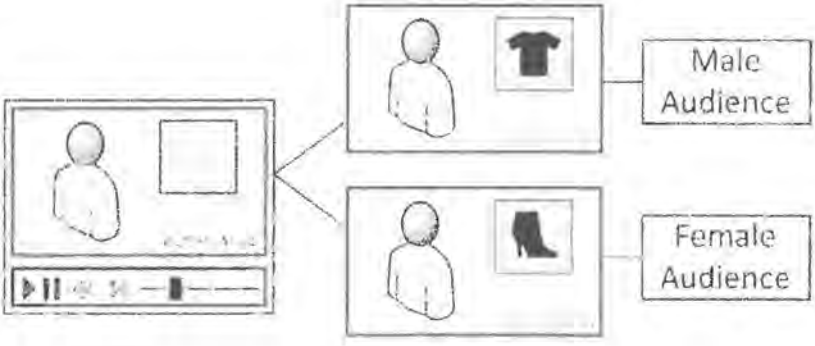


Figure 3.5.7 *DCO is available for Instream. Even with video that has already been filmed or made, the footage can be adapted by inserting a dynamic area into the video that even changes with the light and movement of the footage. This area can be mapped to other images or 3d objects as assets; the outcome is new versions which can be applied to audiences or placements just as with DCO for standard Display.*

Technology permits scalability by allowing the user to create many versions, inserting different, relevant messages into these subtle locations and exporting different versions of the video in doing so. Such technologies are in their infancy and therefore are not altogether cost effective at this advent. On example in

particular: *Impossible Software*; have the ability to plug such a feature into ad serving technology. In this workflow the creative developer would pass the single video (which would be used as the template) into the user interface of the video versioning tool and the trafficker would control the output of the versions into the third party ad serving interface. Subsequently these versions would be assigned to placements and audiences (see Chapter 4).

Beyond VAST and VPAID

Digital Advertising is overly ripe with acronyms and confusingly a VPAID ad can be served through a VAST tag. The important thing to remember is that VAST is a tagging standard whereas VPAID is more like a format or a standardised way to describe the functionality of an interactive video ad. As of publication the so-called standardisation effort has given birth to a brood of numbers and letters that without explanation have little logical labelling:

- VAST 1.0
- VAST 2.0
- VAST 3.0
- VPAID 1.0
- VPAID 2.0

Without getting too technical each is a significant improvement on the other but since there are several standards for video in circulation, third party ad servers need to certify that ads will display and function when run with premium Publishers and broadcasters. This is both to check that the claims that a Publisher can support both VAST and VPAID and if so to see which version they are supporting. In turn the ad server needs to ensure it supports the output to the corresponding standards.

Just as Rich Media Publishers used to require certification the process remains the same. When a media plan becomes available from the Advertiser direct or the Agency media planner, the list of

Publishers should be supplied to the third party ad serving provider so that, if required, a certifications team can perform certifications. Certifications are only required on new sites and players but where these exist, the timeline for the campaign management and setup process will increase to accommodate a full certification.

Publisher Certification

In Instream ad delivery, the video player environment on the Publisher site that renders the ad tag is not standardised, not everyone is using the same kind of player technology. Therefore each Publisher site needs to be certified to ensure that third party served instream ads are fully functional in terms of delivery, verification and reporting.

Efforts have been put in place by the *IAB* and other interested parties to standardise the delivery of Instream ads industry wide with limited success, a claim of usage of the VAST standard by a Publisher is no guarantee the practice is being undertaken. Unfortunately due to the sheer volume of video player technologies in use, even the third party ad server has extended processes around running video campaigns and media plans must be checked. If new sites and new video players appear on the scene, testing and certification are sensible undertakings.

These processes are common place when a channel is in its infancy (and the mobile channel requires similar certification to ensure that the third party ad server tracking will work. The key is to test ad tags so that there is no discrepancy between ad server and Publisher reporting before the campaign goes live. This way Publishers can optimize ad delivery so that there is no over or under delivery with almost identical reporting as the third party ad server (where the Advertiser or Agency will be monitoring the delivery independently).

The certification process is actually very straight forward and with responsive parties, the process need not take more than a few days. Having said this, in practice, up to three weeks is recommended between the announcement of a new video Publisher on the media plan and the completion of the long term certification of that Publisher.

- Third party ad serving providers have a Publisher certification team in place who, once alerted to the presence of a new site on a plan, make contact with the Publisher verbally and send over test tags in an email to the Publisher contact.

- The Publisher takes these test tags and loads them up onto the Publisher-side ad server, displaying the resulting ads in a test environment.

- The location of the test environment is supplied to the third party ad server certifications team who run some tests to check that the metrics are collecting correctly and that the ad server calls are all taking place at the right time.

- When both Publisher and ad server satisfied with the metrics they are seeing the Publisher is labelled as certified and the Advertiser can go ahead to run activity on this Publisher in this channel.

- Certification usually means that any Advertiser from that point forwards can skip this process to run video on the Publisher site in the tested player.

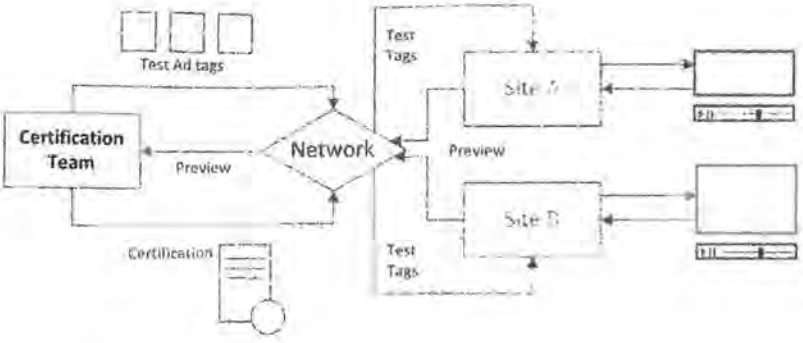


Figure 3.5.8 *Certification for Instream remains a requirement in digital advertising. The certification team at the third party ad serving provider sends test tags to the Publisher direct or the network for the site operations team to implement the test tags in the correct player environment. The preview of this implementation is sent back to the Certification team who assess the preview, awarding certification if the preview passes various tests.*

Certifications can be lengthy processes, especially when the certification reveals discrepancies and errors. For this reason it is recommended for the media planner to supply the list of sites and publishers for the video channel to the third party ad serving provider long before the campaign live date. Owing to the last minute building of media plans, Media Agencies can sometimes find themselves in trouble trying to push video to uncertified sites. It is therefore very important to be aware of the full scope of the certification process for Instream before ever buying the video ad space.

Publisher certification as a process is nothing new. Rich Media was also once an infant channel and used to follow the same process. Eventually all the established publishers and players supporting video certified in full and although re-certifications need to occur when there are big site changes, the certifications all go on quietly in the background without disrupting media plan timelines too much. Rich Media as a channel has now matured

and Instream will follow suit and standardisation improves; ridding this extended campaign process in the coming years.

Trafficking Instream Ads

As mentioned earlier in this section; Instream video ads tend to be simpler to traffic from the third party ad serving interface than Display or Rich media ads which will come as a relief to all traffickers. Linear VAST ads mean the simple vetting of and placement assignment of core flash files (*.mov, *.flv, *.avi are examples and all of which are very common video file types). With VPAID the process is more complex but as an immature channel VPAID in Instream is still rarely made up of complex interactive executions.

Traffickers are encouraged to obtain the highest video quality available from the video producer or Creative Agency. With a higher definition video asset, the ad serving interface can transcode the video into lower quality videos for sites and Publishers that prefer to work with a lower video bit-rate. Ads are created, assigned to placements and placements to sites in the same way as with Display trafficking before the tags are generated and sent to the Publishers. With VPAID, the process of testing interactions in a preview environment is also adopted.

With some ad servers the trafficked ad tags look a little different than Display ad tags when the code packages are opened on the Publisher ad ops side. The ad tags contain links to various qualities of the video to allow the Publisher team to select which version they would prefer to serve to their users. More advanced Publisher ad servers assess the suitability to serve a certain bitrate (usually based on the speed of the internet connection) and serve the video with the bitrate deemed more suitable by the Publisher ad operations team.

With Instream ads trafficked, the tracking feeds down into the third party ad server when the ad is served and the metrics are

recorded (see Chapter 5 - Channel Specific Reporting for details regarding these metrics). As with all other channels, the unified tracking and reporting opens up the possibility to understand how Instream ads add value to campaigns running in other channels. Hopefully such insights justify the higher cost to serve. The consensus remains that this channel provides a more engaging impact upon the user which will encourage better performance of the campaign as a whole.

Channels which can deliver visual creative are hotly sought by big brand marketers who use the impact associated with imagery to draw user attention. With such creativity there is limited call-to-action within the creative except for activities that might encourage better education or awareness about the product or service. Rich Media and Video digital advertising are very well suited to this need and on the periphery sit Display and its maturing cousin; Mobile. The next section divulges the campaign management process for Mobile, disclosing the key concepts behind its gradual refining and why as a channel it is so intangible. Unlike Instream which might be a relatively simple chore for the trafficker, mobile continues to produce headaches for Agency operations teams the world over in the context of third party ad serving.

Chapter 3 - Section 6
Mobile

- The Connectivity Consideration
- Internet Enabled Device Divergence
- Publisher Certification for Mobile
- *Figure 3.6.1*
- Challenges within the browser
    - Responsive Design
    - Whatever happened to Flash?
        - The fall-back Mechanism
    - Trafficking ads for Mobile Browsers
    - QA for mobile browser-based ads
    - *Figure 3.6.2*
- Challenges within the app
    - *Figure 3.6.3*
    - Trafficking ads for In-app ad campaigns
        - Where do Mobile Ad Networks add value?

The Mobile channel presents the greatest upcoming opportunity for Advertisers. The growing volume of internet traffic coming from mobile devices is staggering providing a channel of enhanced reach. Since the ratio of devices to humans is reaching a level playing field, devices have become unique to each human which means the eradication of the assumption that a whole household is accessing the internet from the same place. Further in this section we will see how this resolves in data collection. Beyond improved reach, the features of what were originally called "smartphones" are standardised enough that creative developers can assume that ads can be developed to make the most of these features. Examples include:

- Gestures
- Geo-location
- Accelerometer

This would suggest that the mobile channel should be an exciting area for growth and creativity but it is fraught with challenges that many technology providers, Agencies and Publishers have been accused of not implementing consistent standards as have been created in the other digital channels.

This section begins by looking carefully at what it means to ad serve to "mobile" devices, the challenges with serving to both browser and in-app environments and the process of trafficking to certified mobile properties. Finally we settle on the politics of mobile as a narrowly controlled media channel and the importance of supporting independent tracking mythologies to disclose reliable reporting metrics.

The Connectivity Consideration

Even with the introduction of Smart Phones, internet connectivity and bandwidth remain an issue for delivering high quality creative to mobile devices. For instance, most Advertisers and Media planners will rule out the possibility of running rich media with video formats on mobile Publishers because it puts a heavy drain on the user: the need to download a lot of information to display the ad. To avoid the user suffering a bad experience on a Publisher site, high file size ads have been discouraged, which limits creativity. Where technology is still catching up is finding a way to identify those users using a mobile device on the move versus those that are stationary and are downloading data through a Wi-Fi connection. If connected the Wi-Fi, suddenly bandwidth and data downloading does not need to be restricted but ad servers are not currently providing device level targeting that distinguishes between Wi-Fi and connection via a mobile provider. This feature may never be needed in the ad server itself owing to the improvements in wireless data transmission. 4G

promises internet speeds that will rival Wi-Fi and if data costs come down over the next few years there will not need to be an ad server feature to separate out the users by their bandwidth speeds. For the time being it is advised for creative developers to build the frugal Publisher specifications available, since this makes full consideration for users who might not be able to download large ads.

Internet Enabled Device Divergence

Since there is a whole plethora of internet-enabled devices and the distinction between what is a mobile device and what is not, is a conundrum. In order to separate out media inventory, Publishers began by creating different properties in the hope that the browsing user would head to the property most suited to their device. With this methodology a user arrives at the standard site from their "desktop device", chooses to download the app from their mobile device or uses their in-built mobile web browser to head to the purpose built mobile web site. Larger Publishers such as newspaper publications have taken this a step further and developed applications designed for use on tablets to give a brighter, larger and more engaging experience. Here the distinction is clear, ads must be created for the browser environment and for the "in-app" environment:

- Desktop browser
- Mobile Browser
- In-app for Mobile
- In-app for Tablet

Making the distinction between the user's arrival is an easy task for the Publisher, the browser type can be detected by the JavaScript design of the Publisher site such that mobile users are re-directed to a mobile site while others land on the usual desktop browser page. For apps, the distinction is made when the user first downloads the Publisher's app; choosing the app meant for tablets versus the one meant for other mobile devices.

Publisher Certification for Mobile

As with Instream Video and originally rich media, both apps and browser Publisher properties need to be fully tested and certified before an ad campaign can go live. Without certification, the Advertiser runs the risk that the data collection and actual ad delivery might fail. As a best practice, when a mobile campaign is being booked and inventory purchased, the third party ad server should be alerted so that if required, the certification process can begin. Just as with the other channels, the certifications team at the third party ad serving provider contact the Publisher directly and send over test versions of mobile ad tags. The Publisher implements these on the correct properties and the certifications team check the ads deliver and function and that data is collected in reporting.

Just as with Display inventory, Mobile inventory is partly held by large Publisher Networks that tend to be mobile-specific and certifying through the Networks means that much of the mobile web is Publisher-certified to run third party tags. As with Display it is recommended that should any new Publisher sites or Networks crop up on a media plan, that these details be supplied to the certifications team around three weeks before a campaign start date. Rich Media and DCO ad formats are enabled within most third party ad serving interfaces and open up the options for what kinds of ads can be delivered and tracked at scale but the use of these will extend the period of certification. Thankfully the IAB has started to standardise some core Rich Media formats specifically for mobile which, as with the Rich Media channel for non-mobile devices has been a huge help in eradicating the need for certification almost altogether.

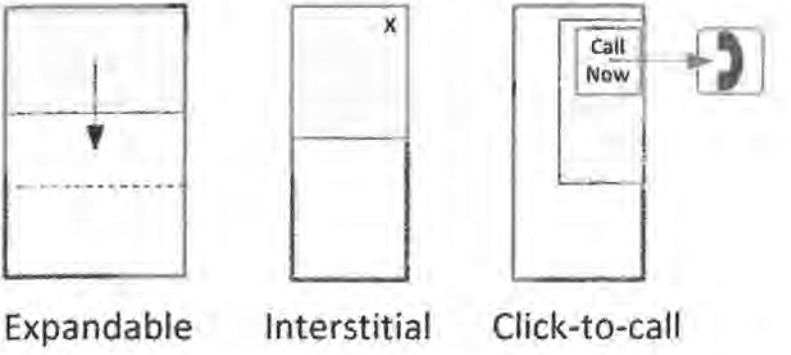


Figure 3.6.1 *Examples of Rich Media units for Mobile campaigns which include expandable ads (simply ads made up of two core panels), interstitial ads (which fill the screen to cover the content much like a floating unit from Rich Media for Display) and the Click-to-call functionality (unique to the mobile space driving the user through to a phone number which connects from the device instantly).*

Beyond certification the mobile landscape is diverse and technically speaking, cumbersome and complex. The distinctions and challenges faced in delivering ads to mobile browsers and in-app environments is great enough to have slowed the spend in mobile channels for media space.

Challenges within the browser

Within the same ad campaign the creative developed for these environments will be consistent in design so as to keep the costs of creative development down. If there were only four possible environments things would be relatively simple but mobile devices differ in screen size significantly. This wasn't always the case, so developing a standard 300x250 meant the same ad could be used on the desktop Publisher site and in the Publisher mobile site. That is until the a user arrives to the mobile site with a small screen and sees just half a 300x250; not an effective execution and

essentially a wasted delivery. When the user arrives at the Publisher site the JavaScript can be used detect the screen-size to request an ad of the most suitable dimensions from the third party ad server by firing off the corresponding ad tag. This still means that the Creative Agency needs to develop many ad sizes and will use the Publisher specifications to determine which ones to build out. The Publisher ad operations team and the trafficker would have a big task on their hands if they had to create hundreds of tags to do this job. Thankfully technology has come to the aid.

Responsive Design

Responsive Design is a new concept but essentially means that with a single ad tag implemented on a Publisher page that the screen size can be adjusted and as the screen enlarges or shrinks, the delivered ad unit changes. The effect is that the same placement appears to deliver a leader-board on a tablet but on a mobile device that spot shows a banner format. The Publisher ad operations team can celebrate, receiving just one tag from the Agency trafficker and yet all users seeing only the ad dimensions most suited to their device. Discouragingly the Creative Agency still needs to design ads to each of those dimensions and although third party ad servers are now offering creative templates in Flash to assist with building responsive ads, there is still likely to be a higher cost to the Advertiser than there would be for developing a single format for a desktop-only campaign.

Whatever happened to Flash?

Flash, the cornerstone of the creative development process is being phased out across mobile devices. But it is not the browsers which are not supporting the plugin required to play the Flash animations, it is Adobe themselves who have

opted to stop offering the plugin for mobile browsers. The challenge to adapt the plugin to an ever divergent technology landscape without profit put the nail in the coffin for Flash. Instead HTML5 presents an opportunity to have the same (if not a better) visual experience than Flash on a mobile device.

The fall-back Mechanism

Within standard display a fall-back mechanism has long been the standard: serve the flash animation, and in instances where Flash is not supported, show a backup image. This method does not serve mobile campaigns very well at all. Since the vast majority of mobile devices do not support flash, users are seeing static images as ads instead of the costly animated creative that was developed for a better user experience. The demand today looks like this:

- If the user's browser can support Flash, show the Flash ad
- If the user's browser cannot support Flash, show the HTML5 ad
- If the user's browser cannot support HTML5, show the backup image

As mobile internet traffic grows further it is likely that the demand from Advertisers will be that the HTML5 version is the one at the top of the food chain. At time of publishing, HTML5 is being treated as an inferior coding environment (where in fact its capabilities are better supported and superior to those of Flash). Although fall-back mechanisms have become integrated into the delivery of ads by third party ad servers, for mobile campaigns, it is still a best practice to look at device delivery reports and to see how many backup images are being delivered by a campaign.

Trafficking ads for Mobile Browsers

The Agency buyer or marketer procures the mobile inventory from the Networks or Publishers and supplies the Media plan to:

a. The third party ad serving provider to check the sites for the purposes of certification
b. The creative developer who builds out the required creative based on the media plan, taking into account the Publisher specifications.
c. The trafficker, who sets up the campaign in the same way as a display campaign

Once the trafficker identifies that the sites being served to are mobile, the ad server grants the ability to attach creative to the ads to support both responsive design and a fall-back mechanism. The trafficker then traffics the ad tags to the Publisher ad operations team who implement them on the mobile site. The total number of tags being kept to a minimum thanks to responsive design.

QA for mobile browser-based ads

To QA the ads the trafficker needs to use a different method than is used in other Display channels since the actual ad needs to be tested on as many mobile device types as possible. In order to check the ad thoroughly and monitor the ad calls, the trafficker can use an HTTP tracer like Fiddler and adjust the proxy settings on the test mobile device so that all the traffic to and from the mobile device runs via Fiddler running on a local desktop based computer. The trafficker can then create a preview URL of the ad inside the third party ad server interface and send to the mobile device from where the URL should be accessed via the mobile browser.

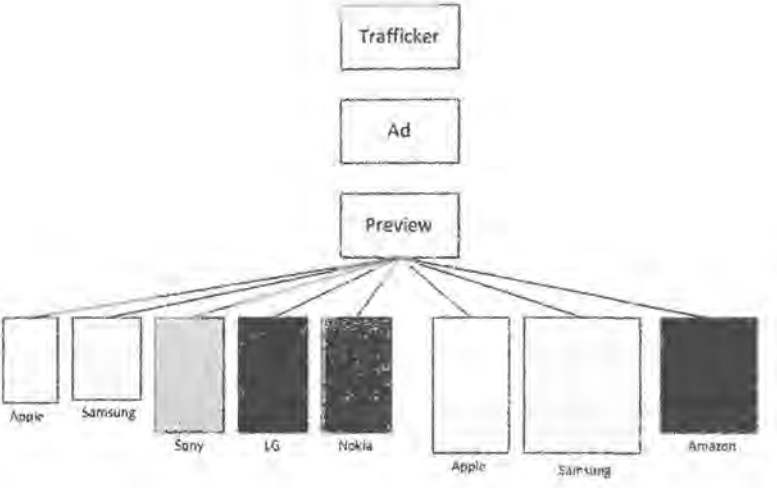


Figure 3.6.2 *Owing to device divergence, a thorough QA process would mean testing a preview link generated by the trafficker on all available mobile devices to check that the ad loads and is at least clickable before the mobile campaign launches. This goes for mobile phones and tablet devices.*

There are some web-based mobile browser emulators available to check ad tags across a range of devices without performing a QA across every device type available but where possible it is recommended for the trafficker to QA the ad on each physical device type.

Challenges within the app

The app space has a growing number of ad formats cropping up and these are predominantly owned by the dominating Mobile Ad Networks. These ad networks control the delivery of ads to the apps making it easy for app developers to monetise their applications by signing up to the networks, deploying a piece of code in the app and having the network serve up ads whenever the app pages load. This has

assisted in supporting wholly free apps, providing developers with a revenue stream.

To third party ad serve an ad into an app means channelling through the network but also delivering the unit and gathering the data into a non-browser environment, so here is where things get a little tricky. Apps are rarely designed to allow for multiple calls into and out of a uniquely developed code structure, so apps much originally be developed using a template that produces a standardised communication channel in and out of the app environment. This template is referred to as an "SDK" (Software Development Kit) which, when supplied by the ad server ensures the basic ad calls will occur sufficiently for accurate counting in ad server reporting. Encouraging Publishers to build applications with such an SDK for the purposes of third party ad serving (and thus greater and fairer monetisation) means working with Publisher-side development teams to have them understand the benefits of putting more work in the development of the application.



Figure 3.6.3 *The SDK circumvents the problem of the black-box environment of the App, facilitating calls into and out of the ad to the ad server or measurement technology.*

Trafficking ads for In-app ad campaigns

Whether a Publisher chooses to work with an SDK or not will still require the third party ad serving provider's certification team to test ads in the application before certifying the environment and approving the app for third party ad delivery on a media plan. Even without the SDK in place, other technologies have emerged on the Publisher network side to facilitate ad tracking called "Ad Enablers". The trafficking process for in-app ad campaigns is the same as in-browser for mobile, where either the direct app or the mobile network are listed as the site in the campaign management workflow in the third party ad serving interface.

To QA how an ad will behave in an app environment is more complex that QA for mobile ads in the browser. The trafficker must work with the Publisher to provide a test page when the actual ad tag is implemented. Very few mobile apps offer emulators which can be used by the trafficker on a mobile device to emulate the in-app environment for the purposes of testing. Without the test the risk is the ad will not function as expected.

Where do Mobile Ad Networks add value?

As mentioned in the Instream Video section, if an Advertiser chooses to run a campaign across a single network then there is no query of de-duplication because users are confined to seeing ads served from one place and tracked by a single system. For this reason Networks try to have as much reach as possible and as in the video world, in the mobile world this is entirely possible. *InMobi* is an example of a high reach mobile network. Third party ad-serving continues to be useful on just a single network to independently verify the delivery and tracking of these ads.

It is clear that the delivery of third party served ads to both the browser and the in-app environments, offer their fair share of challenges. Only certification at its most thorough can come close to the guarantee that the ad will deliver with the support of a fall-back. Data collection for these ads is not quite as standardised and straightforward as for the other display channels and this comes back the device itself (more detail on this in Chapter 5 - Channel Specific Reporting).

As Mobile matures it will take pride of place among channels that are more purpose built to drive specific actions like sales. Channels which drive specific actions like sales are referred to as Direct Response channels. The most significant among these is the Search channel. The next section in this chapter unwraps Paid Search (also known as PPC) which is a much more mature digital channel that is fully third party ad server supported and is seen as the gateway to performance marketing.

Chapter 3 - Section 7
Paid Search and Natural Search

- SEM and PPC
- Creating Campaigns and Keywords in the Engine
    - *Figure 3.7*
    - *Figure 3.7.1*
- Connecting a Bid Management System to the Engine
    - *Figure 3.7.2*
- Why track paid search through a third party ad server?
- How to track Paid Search through a third party ad server?
    - Manual Search Tracking
        - *Figure 3.7.3*
    - Automated Search Tracking
        - *Figure 3.7.4*
- Ad Server tracking integrated with Bid Management
    - *Figure 3.7.5*
- Extending tracked Paid Search to the long tail
    - *Figure 3.7.6*
- Extending tracked Paid Search to Mobile
- Search and other Channels
- What is SEO?
    - *Figure 3.7.7*
- Referral URL to trigger a click
    - *Figure 3.7.8*
- Ad Serving costs for Search

The next areas of this chapter provide reference information for what are called the "Performance Channels". Performance Channels are highly measurable channels marrying reach with effectiveness; specifically designed to drive audiences through to conversions (see Chapter 4). Performance Channels are located

lower down in the strategic marketing funnel because they are explored and used by consumers that have predominantly come closer to finalising the purchase decision. This can be compared to the display channels which often better serve the needs of the Advertiser further up the funnel to create brand and product awareness.

Search is effective when users have already built up an awareness of a product or brand through more visual channels such as display, as they use the search engine to research the internet for more information. In this sense Search engines are the navigation hub of the web experience and as such are valuable places for the Advertiser to attract attention.

Search can be divided into two channels Paid Search and Natural Search. The two channels are almost never managed by the same team but both channels can technically be tracked by a third party ad server. It is worth noting that for the Search channels, the ad server is not considered to "serve" anything since there is no image displayed but rather can attend to a portion of the user click path for tracking. The ad server's position is still wholly relevant because internet users that might take actions with any other channel, may also perform an action in a Search channel and capturing this pathway is essential for conversion and attribution reporting (see Chapter 4).

This section begins by looking at how paid search is used as a channel by the Advertiser and how the third party ad server connects to the search engines to seamlessly integrate for tracking purposes. This process has evolved over the years and has taken a big workload off of the trafficker's shoulders. We will touch on how the shady past shapes the process today and why the future of campaign management for paid search looks so strong. This section will also touch on bid management solutions; which assist the core ad search programs to scale the implementation of keywords. More controversially this section

will end by explaining how natural search can be tracked as a channel in its own right and how the setting up of this channel from a tracking perspective must be carefully considered so that more directly paid-for channels do not appear to lose their effectiveness.

SEM and PPC

Paid Search is often given the acronym "SEM" which stands for Search Engine Marketing, however, Paid search channel is also commonly referred to industry-wide as PPC, which has the potential to confuse those new to the world of search since PPC actually means Pay Per Click. It is true to say that the paid search channel is predominantly funded through a PPC or click-based pay model, as opposed to the CPM dominated model for the Display channels, but it might not be the best name for this channel as a whole.

Defining Paid Search is fairly straightforward but this section will not dwell too much on its intricacies beyond the intermediate educational requirements of the Advertiser to third party track its activity, since there are a great horde of how-to textbooks already written on the topic.

Paid search is essentially the sponsored search listings that appear on the search results page in the search engine. These traditionally appear along the top (just under the search bar) and down the right side of the page. They are usually presented on an off-colour background and the ads presented usually match the keyword typed in, in some way. The search engine algorithm for the paid search listings is meant to ensure that Advertisers that have paid more for their keywords to display will appear higher up the results list or more frequently across more users search results pages.

When referring to the paying of keywords it is useful for us to loop back to the term PPC. The search keywords themselves cost

the Advertiser money only when there are clicked on (the delivery of the sponsored listing itself is essentially free). Advertisers arrive into an auction to offer up the maximum price they would pay should a user click on their ad. Winning the auction will trigger a charge in the Advertiser's account.

In its most basic form, the Advertiser is permitted to enter into this auction by creating an account for buying keywords with the Search Engine. For instance *Google* offers a program called *Adwords*, *Microsoft's Bing* offers *Adcenter*.

Creating Campaigns and Keywords in the Engine
Paid Search activity is usually managed by a specialist search team which belongs to the marketing team on the Advertiser side directly or sits in the Media Agency (separated from the Display operations team). Some Media Agencies have carved off a decent portion of their staff to create new search specialist Agencies (a good example is *Performics*; an SEM Agency under the *Publicis Group* umbrella). These teams access the SEM interface belonging to the engine by navigating to their site and logging in.

The first step to setting up a search campaign for the first time is adding bank details and funding the search account. Once funds have been deposited, the team can begin by creating or uploading keywords. Keywords are basically anything that a user might type into the search engine to find something. Search engine marketers are encouraged to create as many keywords as they can think of and many user specialist tools and programs to generate keywords in their thousands and millions.

Once the keywords are determined they need to be coupled with a short description and the URL where the user should be directed to once they click on the ad, essentially like with ad serving; a click-through URL. Keywords coupled with the description, title and clickthrough URL make up the search ads themselves. To help manage these ads they are sorted into buckets in the SEM

interface: called "Ad Groups" which in turn are clumped together to form campaigns. The most common way to mass produce millions of search ads involves using spread-sheets on a local machine to couple the generated keywords with their component parts and the name of the ad group to which they belong. The sheet is then loaded up into the interface underneath the correct campaign.

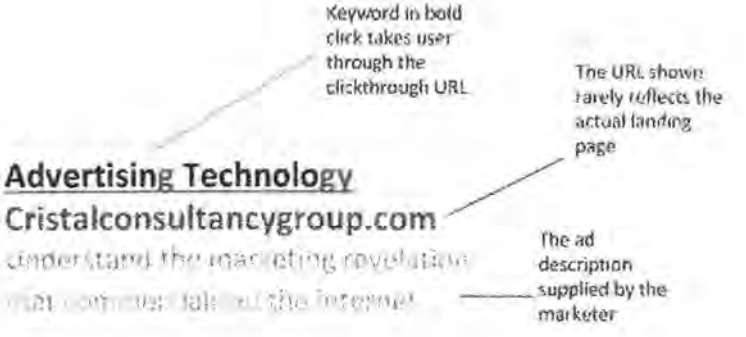


Figure 3.7 *Components of the now standardised Search Ad include the Keyword as the header underlined in bold followed by a URL (which does not have to reflect the true landing page or the clickthrough URL). Both are accompanied by a short description.*

Campaigns in the SEM interface can be defined at the marketer's discretion, many marketers are following a convention of having a new campaign for each brand, product and marketing campaign. The campaigns listed in the interface may not need to all be active at the same time, so it is simple to activate and deactivate the activity. Also alongside each keyword the marketer needs to determine how much they are willing to pay for the keywords so the maximum bid price is entered. The spread-sheet as a whole is then uploaded to the interface and the campaign can be put live.

Inside the same SEM interface in the engine, it is possible to get reports to show which keywords were clicked on and how much

was eventually paid. Despite its simplicity there is a whole range of options available to enhance the performance of the search campaigns, from adjusted messaging at different times of day to utilising retargeting strategies (see Chapter 7). What is actually happening to the user when they click on a sponsored search result in the search engine, is that their click activity sends a count to the SEM reporting server so that the information about the click such as the keyword clicked on and the timestamp of the activity is stored just as the user arrives at the landing page destination.



Figure 3.7.1 *The click-path that the user is taken on as a result of the action of searching. In this example the user navigates in their browser to the search engine homepage (in this instance Google) where their search term is carried to the search results page. Here the SEM Search Ad is displayed. The user clicks the ad (which is*

labelled as a sponsored listing) and appears to land directly on the homepage. The user is in fact redirected through the Google Adwords reporting server so that the click is registered and the Advertiser is charged for the activity.

Connecting a Bid Management System to the Engine

The search engines have developed their own APIs to allow other systems to connect and control this keyword and campaign management environment, at a distance. What these systems have that the search engine SEM interfaces do not, is more complex ways of deciding what keywords to pay more for over others. In the context of the auction model in which PPC is determined requires that at scale, prices be set according to the value that each keyword is bringing. Since a multitude of factors determine this, these external systems are commonly referred to as "Bid Management Systems" and pull on various data sources to make decisions.

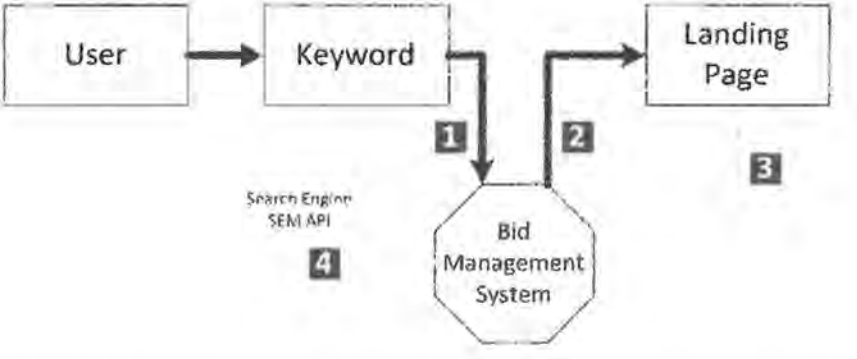


Figure 3.7.2 *The user would typically input a search term into the search engine such as Google and the ad clicked on (represented here by a keyword) would take them to the landing page (redirected through the Adwords reporting server). With a bid manager involved the user's click-path will still route via the Adwords reporting server but will also route via the reporting*

server for the Bid Management System. This is represented by (1) and (2). Information about the user's journey is also collected by the Bid Manager's technology (also embedded on the Advertiser's landing page) and this is also fed in the Bid Management System's reporting server. With information about clicks and conversions (see Chapter 4) the Bid Manager can update the bid price on the keyword via the Search Engine's SEM API.

Bid Managers track the user as they click on the keyword to determine how far that user goes down the purchase funnel and adjusted bids for keywords based on this as a starting point. To this extent, when the bid manager is setup on the Advertiser's SEM account, the bid manager connects to the search engine via the API and adds a "code wrapper" around the click -through URL of the keywords listed there. The wrapper is portion of code that sits either side of the click-through URL and when the user clicks the link, the bid manager is now tracking the user independent of the search engine.

Despite having their own SEM interface; *Adwords*, *Google* has its own Bid Management tool which makes up part of the *DoubleClick* suite; the latest iteration at time of writing is called *DS3* which evolved out of the legacy tool *DARTSearch* which was one of the first ad-server integrated bid management tools. Other popular and powerful players in the space are *Marin*, *Kenshoo* and *IgnitionOne*.

> Bid management tools have begun evolving in recent years so that they take on new data from the Advertiser side and begin acting a little more like a DSP or a DMP but leading with bidding on the search channel (see Chapter 8 - Programmatic). The difference here is that Bid Managers are built for enormous scale for the handling and processing of millions of keywords where bidders in the Display space are working with far fewer ads.

Having looked into how campaign management works for search marketing campaigns; both smaller scale inside the SE (Search Engine) interfaces and at a larger scale with bid management systems there remains an important space for the independent tracking and counting of the data by the third party ad server.

Why track paid search through a third party ad server?
Owing to the scale and dynamics of Paid Search, it is considered to be a wholly separate channel from Display advertising. This means that in order to track users that:

1) clicked on paid search
2) were influenced to search by display ads

..there must be a common technology tracking both channels at a single point. The regular use cases for using a third party ad server still stand; such as the case to have an independent tracking technology, kept separate from the Publisher/Search Engine provider. It also becomes possible to compare and contrast the costs of paid search across multiple products, brands and countries to other channels in a single system of record with the same counting methodology.

How to track Paid Search through a third party ad server?
To the trafficker working inside the third party ad server interface, the search engines can be classed as Publishers and the keywords as the ads or placements. The trafficker navigates to the Advertiser account in question in the ad server hierarchy and there are two methods to then connect the third party ad server from there.

Manual Search Tracking
Manual Search tracking is how things used to be done in the world of ad serving for paid search. This method is now reserved for very rare instances in which the Search Engine

API fails or for connecting to smaller more obscure search engines. The trafficker undertakes the following steps:

1. Sets up a new campaign in the third party ad server interface with the correct naming convention (that might mean the advertiser is naming different campaigns according to which channel they belong to) or accesses an existing campaign that contains activity for all channels
2. Generates a tracking pixel with single click tracker for each keyword in the engine inside the campaign
3. Inserts the same landing page URL as can be found alongside the existing keywords in the engine into the click trackers.
4. Traffics all of the click trackers, so that the ad tag code for the clicks becomes available
5. Downloads the keyword list from the Search Engine SEM interface for the campaign in question
6. Replaces the click-through URL column in the keyword list spread sheet with the matching code for the click from the ad server
7. Uploads the completed spread-sheet back to the Search Engine SEM interface and puts the campaign live from there

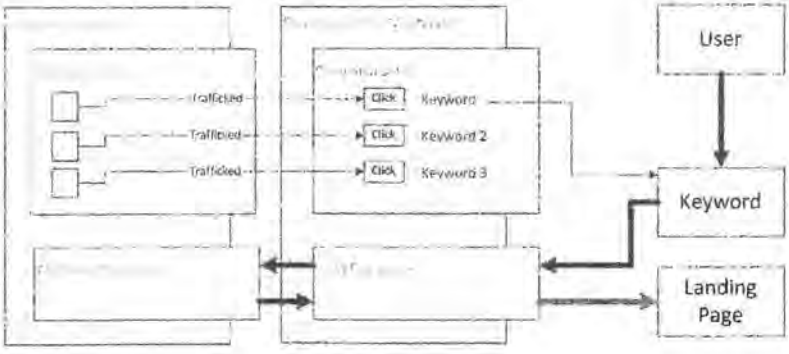


Figure 3.7.3 *Manual Search Tracking involves a heavy operational burden on the trafficker. A tracking placement is created for each keyword to match the keywords in the engine. The click redirects are then trafficked to the Media Agency who replace the click-throughs on each of the keywords with those generated by the ad server. When the user clicks the ad for the keyword in question, the user click path is routed through the Search Engine SEM reporting (such as Adwords) and the third party ad server reporting before landing the user on the landing page.*

This manual process is enormously time consuming and problems occur when trying to differentiate the click trackers in ad server reporting from other kinds of ad servers if a naming convention has not been set up to divulge the channel type. Thankfully the Search Engines themselves came to the aid of marketers whose total number of keywords was becoming unmanageable with this system.

Automated Search Tracking

The search engine API allows the third party ad server to connect directly to the engine and write the click-through URLs to the keywords themselves directly. The trafficker then:

1. Logs into the search engine from a slightly different screen inside the third party ad server but still sitting within the same Advertiser account. This login will be the same as the one used by the search marketers to access the SEM interface in the search engine directly.

2. Sees the campaigns as they were created in the search engine and selects the right one to track all keywords in that campaign so that the tracking is fed into the ad server under this account.

The third party ad server will then automatically 'wrap' the landing page URLs in the engine with ad server click tracking with no additional effort required from the trafficker.

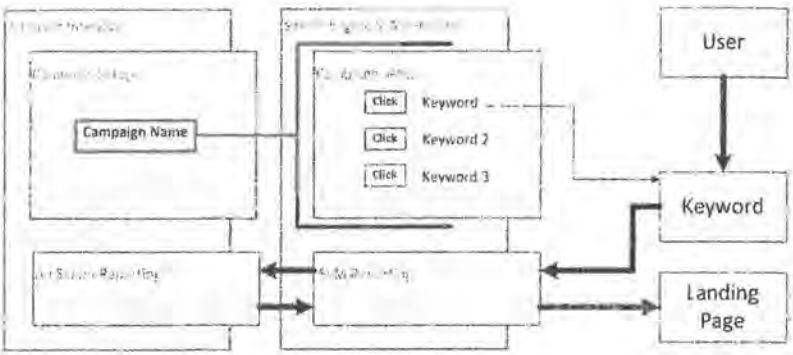


Figure 3.7.4 *The automatic search tracking option presented by some third party ad servers is much more favourable operationally speaking. Here the trafficker only has click "track" by selecting the right campaign name in the third party ad server and the corresponding name in the Search Engine SEM interface has all its keywords wrapped.*

Once the keywords are being tracked the ad server connects to the engine at least once per day via the API to request reporting information which is pulled down into ad server analytics so that the trafficker can pull reports about the cost of the keywords and their performance. This means the Advertiser or Agency does not need to log into the various engines to extract the results but can get them all combined in ad server analytics. More information on this can be found in Chapter 5 - Channel Specific Reporting.

Ad Server tracking integrated with Bid Management

Bid management systems can be used in conjunction with third party ad server reporting and at a deeper level the two technologies can share data independent of the search engine itself. Operating bid management systems and ad-serving is easy for those ad servers that claim to have a basic level of integration. Essentially both bidder and ad server work separately but harmoniously in that their code wrappers both wrap the keyword in the engine via the Search engine API and the two code wrappers do not disrupt the different systems from tracking and reporting separately. To achieve this state the ad server needs to have worked with the bid management provider in a process of certification. In some instances the wrappers disrupt one another unless one wrapper is added after the other, this tends to be the case with lesser known bid management providers and less user-utilised search engines.

Sharing data between ad servers and bid managers is a relatively new idea compared to the longer history of search marketing. In this instance the ad server is collecting a vast amount of data about the Advertiser's ad campaigns outside of the search keyword but can tie all this information back to it. With the Advertiser's approval, the ad server can pass data that might show if some keywords are more or less effective depending when the user saw or acted upon them. In turn the bidder can improve its bidding algorithms through a system of learning.

> An advanced example of such an integration would include the passing of E2C (exposure to conversion) or P2C (pathway to conversion) data from the ad server to the bidder using a daily feed (See Chapter 5 - Cookie Level data). Using this method allows the bid manager to apply an "Attribution Model" to users featured in the data to adjust the bid price for certain keywords (more on Attribution Modelling in Chapter 4).

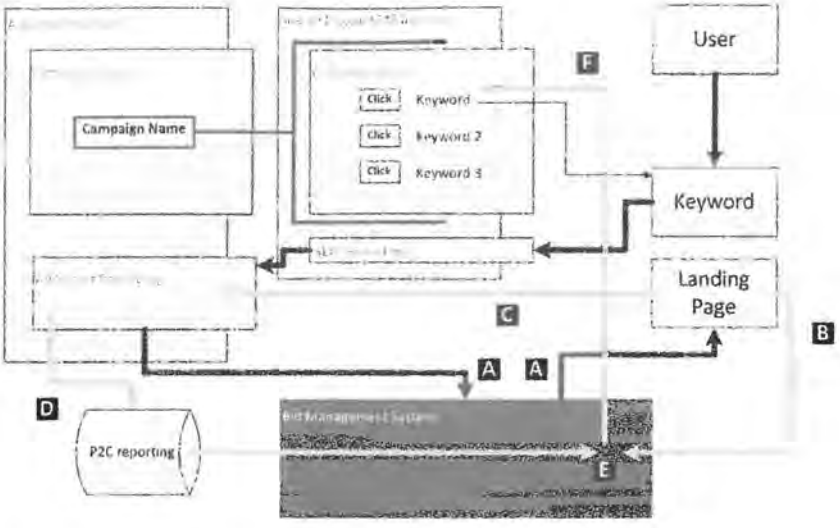


Figure 3.7.5 *Introducing integrated bid management into the equation allows the bids on the keywords to be much more intelligent and the prices to pay to be far better informed. The user click-path now triggers counts in SEM reporting (such as Adwords). Ad Server reporting and Bid Management reporting at point (A). The user lands on the landing page triggering the Bid manager's landing page tag (B) demonstrating that the keyword can be given higher value since it drove a site-side action. The landing page also triggers a conversion tag (C) - more info on this in Chapter 4. The user data about the click and the conversion as well as user data about other engagements across other media on other channels for the clicking users is passed into (D) "P2C Reporting" (see Chapter 5) and the data is collected in the Bid Management system at (E). Inside the Bid Management System, the marketer is free to apply an attribution model (See Chapter 4) to the keyword information to establish its value before adjusting the bid price at (F).*

Extending tracked Paid Search to the long tail

Google Adwords product has a sister toolset called *Google Adsense*. *Adsense* allows small Publishers to attract CPC (Cost per click) revenues through Advertising, which otherwise would be completely unattainable since the amount of user traffic to such sites remains very small. Niche Publisher sites have very specific content that, despite having a small amount of traffic might have a perfectly suited audience for a "contextually targeted" ad.

Adsense works by having the Publisher download *Adsense* code from *Google*'s site (after signing up for an account and inserting bank details) and sees the Publisher embed the code on any page of their site. Users then arrive at the site and the *Adsense* code reads the content of the page. This content is then checked against the *Adwords* databases to see if there are any matching keywords which it then displays to the user. These are *Adwords* paid search ads, appearing in the *Adsense* spot on the page. Since *Adsense* draws from the *Adwords* database itself, *Adsense* uses the same click-through URLs as adwords, if these keywords are tracked in the third party ad server they will also be tracked by *Adsense*.

> Advertiser's should have the option of separating out the traffic coming from *Adsense* supported Publisher sites versus those coming from the Search Engine itself by only applying certain SEM campaigns in the SEM interface to either *Adwords* or *Adsense* since the traffic quality might be different. It is worth the Advertiser checking with the Search engine directly to see if campaigns be separated out. If they can, the trafficker should also separate out these campaigns in the ad server so that it is clear in the ad server reporting.

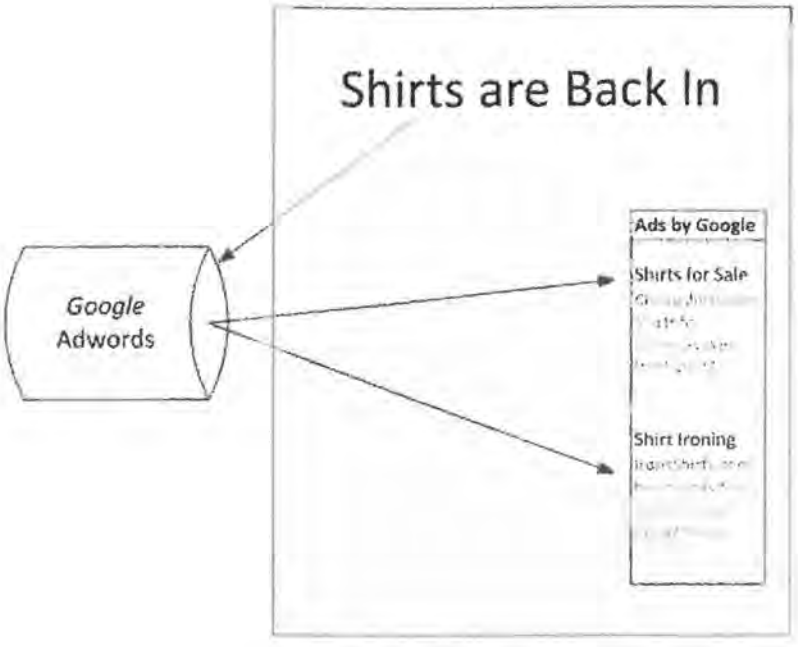


Figure 3.7.6 *Google Adsense is powered by the Google Adwords engine. When the user loads the Publisher page, the Adsense code collects the contextual information on the page and passes it into the Adwords engine before the relevant ads are displayed. This diagram is simplified however Adsense will rarely scrape the page in real time. Rather the Google spiders will have cached the page in the Adwords engine in the past week to associate the URL with the relevant context, then Adsense just sends the URL to engine which checks ads associated with that contextual information which are then served out. The click-through-URLs will still contain the pathway as used in the third party ad server to count the clicks there.*

Extending tracked Paid Search to Mobile

Just as with separating out *Adwords* campaigns from *Adsense* campaigns, the Search Engine should discern those campaigns delivered to Mobile from those delivered to desktop devices. Mobile devices offer the geo location facility and as such, the

search keywords are likely to be more relevant to user and therefore more valuable to the Advertiser. When bidding for keywords in the SEM interface or through the bid manager, the geo location relevant ads are likely to cost the Advertiser more if the resulting activity is likely to bring in more revenue for the Advertiser. If the mobile search activity draws from the same database of keywords and click through URLs that the marketer sets up in the SEM interface, in the bid manager and in the third party ad server then it will be tracked in third party ad serving. Again, to break this out into the correct kind of reporting, the campaigns should be sufficiently labelled in the Search Engine SEM interface and created as separate campaigns in the ad server under the same Advertiser account.

> *Google* have utilised their technical ability to allow Advertisers of all sizes to take advantage of advertising to a bulk number of Publishers through "GDN" the *Google Display Network*. This is not search activity but is more closely aligned with Affiliate advertising covered off in the next section.

Search and other Channels

Tracking and counting paid search provides a wealth of data within the third party ad server that can be used to create inventive strategies to improve the optimization and serving of creative across other channels. For instance, if users are searching on specific keywords, those can be paid for through SEM and the traffic can be driven to a specific Advertiser landing page which may house a retargeting pixel (See Chapter 7) which could be used to power a follow up DCO campaign if the user does not go on to buy the product in question.

Paid Search is often cited as being one of the most effective marketing channels for driving performance activities. But from the Advertiser's perspective, any channel that is effectively driving user web traffic through to the Advertiser site should be funded

by the marketing team so as to make those channels even more effective. Owing to the way Search Engines work, paying for an expensive website that is Search Engine optimized can be an effective use of this budget and SEO (or "Natural Search") therefore is transformed into a channel that needs to be tracked alongside the other channels. It is clear to see how this channel differs from the others, there is no procurement of advertising space, yet the Advertiser's site appears on the search results page. It is important to touch briefly on how this works.

What is SEO?

SEO (Search Engine Optimization) describes the on-going development of a website for it to be picked up naturally by the way a search engines work such that the pages are displayed in the standard search listings. These "natural search" listings appear in the centre of the search results page and trail down the page onto multiple pages. The copy and links which make up these results are generated by the search engine. As the results page loads information comes from the search engine database which is kept updated by the search engine's web crawling technology which are called "Spiders". Spiders are continuously crawling the web in accordance with the natural search algorithm which ranks which pages are the most relevant to user's search term and returns to the database the right information.

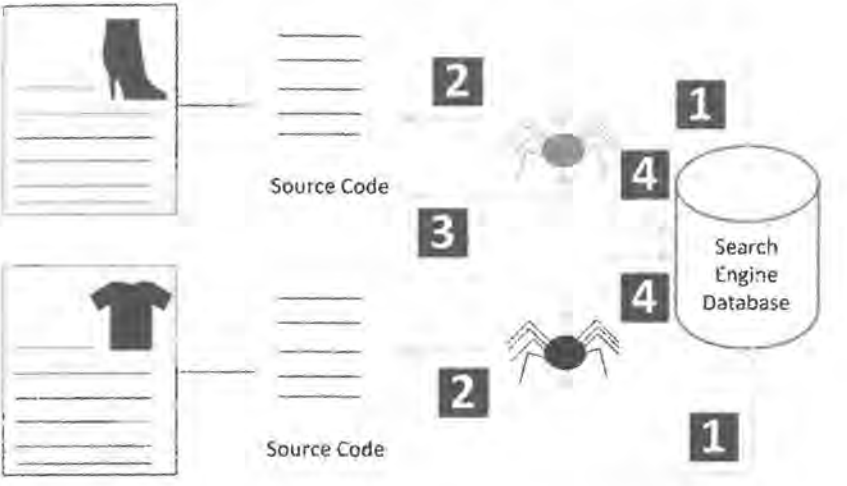


Figure 3.7.7 *The spiders or web crawlers from the Search Engine (1), crawl through the web's URLs to reach websites (2). The crawlers read the source code of the page (3) and pass the required information about that URL back to the Search Engine Database (4). The information is stored as a cache in the Search Engine database. Crawlers or Spiders repeat the process of crawling pages more often depending on how popular they are which can be determined by the number of other domains pointing sending users traffic to the URL.*

Publishers and website owners work very hard to try to get their page seen at the top of these results and because these results are unpaid, the cost to the advertising effort is indirect. As stated, it would not be true to say these results are free since it is the sweat and tears of the website development team that ensure a high "page rank" in the engine. To justify the cost of the man hours put into developing a great SEO strategy, the SEO results need to be tracked to establish the effectiveness of the strategy.

Site side analytics tools such as *Adobe's Omniture* or *Google's Google Analytics* are very effective at detecting the incoming site "traffic" from these engines and understanding where it came

from to quantify the effect of the SEO efforts. Having said this site side analytics tools can rarely tie together a user's behaviour with the natural search keywords alongside other paid for channels such as Display and Instream Video. Also such results are better verified by an independent third party so as not to allow for favour with the Publisher owning the media that is driving the traffic in the first place. The third party ad servers were never originally designed to track this kind of activity but this need is in increasing demand from Advertisers and Media Agencies.

Referral URL to trigger a click

As a reminder from Chapter 2 - Advertisers and Marketers, the user that arrives at the Advertiser site leaves a trail behind them for the last action, giving away how they arrived at the site called a "referral URL". The referral URL gives the Site Side Analytics tool an identifier from which the search term (that the user entered into the engine to get to the site) can be obtained. Site Side Analytics tools do this by using JavaScript. This same JavaScript method can be written out by the website developer, separate to the analytics tool and be used to trigger a click tracker. Suddenly the user receives an ad server cookie where they would not have got one from the natural search result itself. This makes it look in ad server reporting that the user clicked on this click tracker to get to the site rather in that the site arrival triggered the click tracker. The result is that the SEO channel is tracked by this simple user action. Such JavaScript code should deployed on all possible landing pages of the Advertiser site that a user might be able to access from a search engine.

Behind the scenes in the third party ad server interface the setup is a little more complex than creating a tracking placement and using the click tracker portion to mount into the script on the page. The trafficker has two options:

1. Either the script can be designed so that each possible keyword triggers a different click tracker. This would mean

the trafficker is given a list of the keywords by the Advertiser website management team (this list can be exported from the site side analytics tool). The trafficker would then create one tracking placement for each keyword and the click portion of the tags be trafficked and sent back to the website management team. In turn this team would tweak the JavaScript on the page so that when the user arrives on the page, the search term that informed their entry point to the site be captured and the corresponding click-through URL be triggered. Reporting for this option is very simple since the trafficker just pulls a standard campaign report and the keywords are already pre-defined in the report. The difficulty with this method is that a JavaScript would not be able to house too many keywords nor would the trafficker want to create or manage too many tags to support it.

2. ..or the script can be designed to pass the exact keyword search term from the referrer URL into a dynamic token in the third party click tracker. This would mean the trafficker only creates one ad in the ad server and having the website management team implementing the click portion on the Advertiser site inside the purpose-built JavaScript. The JavaScript would be tweaked so that as the user arrives to the site the click is fired and the keyword portion in the referrer URL is dynamically inserted into the click call back into ad server reporting. The trafficker would then need to ensure that when reporting on the clicks from this ad in the campaign reports that the contents of the dynamic token is displayed and aggregated so that the most popular keyword entry points can be ascertained. This kind of reporting can be harder to set up.

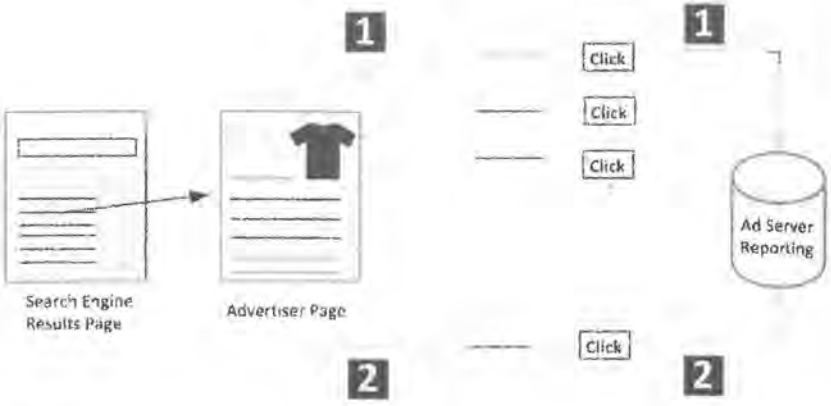


Figure 3.7.8 *The user types a search term into the search engine and is taken to a search results page. On this page the natural search results are listed running down the page which match to relevant sites. The user clicks on one of these results and is taken to the Advertiser site. One the landing page is embedded a script that either fires off the corresponding ad server click tracker for the search term (1) or inserts the search term into a dynamic token in the single click tracker option available to the Advertiser (2). In either instance the information is supplied to ad server reporting.*

If these setups are followed these clicks will be treated like clicks in the Display channel or clicks in the Paid Search channel and this will create a situation where conversions (see Chapter 4) on SEO under the last click attribution model (see Chapter 4) will take the full credit for the conversion and actual paid-for media will be eclipsed by this activity. Although this may sound complex: what it really means is that natural search will suddenly appear to be more effective than other media channels tracked in the ad server: so much so that it might convince the marketer to stop paying for the paid-for channels altogether. This is something of an illusion because if the marketer stops paying for media in paid for channels, the site side analytics tools will report a complete

drop off in site traffic altogether and conversions will also fall short (see Chapter 4). Natural search results are very commonly used but they should be thought of as one of the influencers in the user's pathway to their transaction or journey on or to the Advertiser site rather than the only influence. For this reason the natural search tracking activity needs to be downplayed in ad server reporting (see Chapter 5 - Channel Specific Reporting).

Ad Serving costs for Search

third party ad servers usually charge Advertisers and Agencies based on impression volumes for ads served. Since no image is being delivered for the Search channels, the provider will normally charge based on clicks only for the click trackers.

This introduction into the world of paid search and natural search exposes in greater detail the importance and relevance of third party ad server reporting and conversion calculation.

The next section continues with the theme of performance marketing, looking at what happens when Publishers begin to realise that the best way to attain user attention is to utilise the Search Engines themselves. It may be true to say that such Publishers value their ad space over their content and as such transform their sites into conduits to channel user traffic through to Advertiser sites. These Publisher websites are called "Affiliates".