

March 2017

HPE Vertica Empowers Adtech Industry to Deliver Data-driven Conversions

Stratecast Analysis by
Jeff Cotrupe



HPE Vertica Empowers Adtech Industry to Deliver Data-driven Conversions

Introduction

John Wanamaker, the US merchant considered an early proponent of advertising, famously said: “Half my advertising is wasted, but I can never tell which half.” Humor aside, in today’s ultra-competitive markets, advertisers need precision, in real time, when it comes to operating and quantifying marketing campaigns. Failure to achieve this can cause serious, and quite un-funny, financial pain for brands, marketers, and agencies.

This need for precision, and perhaps a more urgently-felt need than ever before for rapid, and quantifiable, results is putting the \$600 billion global advertising market through some fundamental changes. With the users who are the targets of ad campaigns more mobile than ever, and brands highly cost-conscious, today’s digital advertising industry must respond by becoming more mobile and measurable. Brands, agencies, and advertising ecosystem providers must find the right mix to reach consumers with the right messaging at the right time in the right context. Beyond that, they must close the deal: they have to (verifiably) drive purchases and other conversions sought by the brands who pay the bills.

Although that sounds extremely challenging, there is plenty of good news here. The challenges the ad industry is experiencing have created broad areas of opportunity for Big Data and analytics (BDA) to make a difference—and we can now confidently report that BDA is up to the challenge. Stratecast analyzes and assesses the total BDA market, and relevant slices thereof, such as by vertical and by region, but another way we segment BDA is by functional categories, as shown in Figure 1.

Figure 1: Stratecast Segmentation of BDA Market by Functional Categories



Source: Stratecast

One of these categories, Mobile, Retail, and Location Analytics (MRLA), is focused on using BDA to help brands engage with customers and prospects anywhere, anytime, and a big part of that is analytics used to drive ad campaigns. MRLA accounted for 17.1%, or \$8.28 billion, of a total BDA

market that reached \$48.4 billion, in 2016. This amount of investment reflects the importance of leveraging data to reach one's customers and prospects, and nowhere is that truer than in adtech.

This white paper addresses the challenges the industry is facing, and analyzes why BDA, and more specifically, analytic insights at the speed of business, are now essential elements in advertising technologies (adtech): empowering people (and machines) to instantly make optimal advertising decisions. The paper then presents examples illustrating how this is helping brands and providers achieve tangible results in their businesses.

Industry Challenges, and Roadblocks to Success

Not that long ago, an advertising executive looking to place an ad in a given network would place a call with instructions, such as, "Increase the bid by 15%." Today, that manual process is impossibly slow and does not scale to the millions of transactions necessary for a busy agency, or corporate in-house marketing team, to manage all ad placements across all platforms.

Ad network providers today need to collect and instantly analyze massive volumes of streaming data, including a myriad of combinations of keywords, user types, device types, and user locations, and match ads with available slots. The network presents the most relevant ads to targeted recipients, then reports activity—meaning conversions, or lack thereof—back to the brand/advertiser.

Once ads are placed, the experience for those running digital campaigns can be like releasing toy boats into a fast-moving stream, "watching from shore" as they hurtle downstream at the mercy of the current, and hoping for the best. They have no way of knowing whether campaigns are outperforming or underperforming until well "downstream," days or even weeks later. This has a number of negative effects, including:

- Wasted advertising expenditures, as the brand continues spending on campaigns that are underperforming—and, just as importantly, fails to reallocate ad spend to campaigns that are driving bankable results.
- Missed market windows and lost revenues, as misallocated campaign spend misses buyers with serious intent and brands lose potential sales.

This is especially crucial for the many providers now operating on a shared risk/shared success business model with advertisers: provider success depends largely or entirely on results. Search partner ad networks are a good example of this: their work is only valuable to advertisers if they are successfully creating conversions.

Bottom line: adtech today moves at incredible speed. And few in the ad ecosystem are up to speed.

Bottom line: adtech today moves at incredible speed. And few in the ad ecosystem are up to speed.

Six Degrees of Specialization—and All Require Near-Real-time Insights

The big opportunity in adtech is to meet the challenges of a millisecond business by effectively harnessing BDA to obtain rapid analytic insights; to leverage those insights to optimize campaigns and offers, to maximize take rates; and to optimize costs through efficiency. Those results can only occur by putting rapid analytic insights to work across the six areas of expertise and specialization comprising the digital display ad ecosystem, as shown in Figure 2.

Figure 2: Six Areas of Digital Ad Ecosystem and Why Each Needs Near-Real Time Insights

Area	Where and Why Near-Real Time Insights Are Needed
1. Content Creation	<ul style="list-style-type: none"> • Focused on developing brand content, and generating and leveraging user-generated content (primarily social media) • Social media moves fast, and brand content often needs to change on the fly; rapid insights are essential to continuously refresh and optimize content
2. Media Planning and Buying	<ul style="list-style-type: none"> • Focused on continuously choosing optimal ad placement venues; some have the resources and expertise to do this in-house; others use agencies • BDA helps make go/no-go decision on handling in-house versus agency • Near-real-time information is required to sift through constantly-changing ad inventory to make the right decisions
3. Presentation of Ad Options	<p>Vendors must continuously present brands with the best ad options, gathering data in real time from these sources:</p> <ul style="list-style-type: none"> • Ad networks • Digital media exchanges • Full-service, demand-side platforms (DSPs) • Agency trading desks (ATDs) • Sell-side platforms (SSPs)
4. Campaign Execution	<p>This is where networks and providers need to act on near-real-time information to get the ads placed, via:</p> <ul style="list-style-type: none"> • Ad servers • Site-serving systems • Third-party ad placement arrangements • Self-serve DSPs <p>Once ads are placed, this is also where ad performance is measured, and where campaigns, messaging, ad venues, and other variables are adjusted to deliver the desired results</p>
5. Media Enhancement	<p>Requires near-real-time data (and knowledge of best practices) to enrich creative content and thus more precisely target audiences via:</p> <ul style="list-style-type: none"> • Rich media • Creative optimization • Data management platforms • Tag management • Ad verification and privacy
6. Business Intelligence (BI)	<p>A flow of accurate, near-real-time insights is crucial in digital advertising to:</p> <ul style="list-style-type: none"> • Effectively leverage data across all areas • Manage each line of business (LOB)

Sources: Stratecast, Interactive Advertising Bureau

In short, those in the ad ecosystem need to have the freshest possible data to create content; decide where to place ads; get those ads placed, and quantify their performance; enhance ad vehicles or messages to more precisely target the desired audiences; and apply analytic insights across the business.

A key decision point is whether to perform the core function of media planning and buying in-house versus having an agency do it, and BDA can play a major role in helping brands make that decision—as well as helping agencies show the cost/benefit justification for why brands should work with them for that important function.



Solution: Crunch Massive Volumes of Data to Obtain Rapid Insights

When considering technologies to help adtech keep pace with the ad business, three criteria are imperative: speed, scale, and query support. An analytical database must be capable of processing up to petabytes of data, and delivering analytic insights in real time. It must feature a scale-out, MPP¹ architecture to process queries in parallel over multiple processors, enabling linear scalability, high availability, and fault tolerance, and to provide a ready migration path for expansion and upgrade. It must support methods for approximate query processing (AQP), which are often the only means of providing rapid response times against massive datasets and are required to handle high-speed data streams. AQP methods such as approximate count distinct and approximate count distinct synopsis offer the option to execute queries against the entire dataset, or to speed response times by executing queries against a compact synopsis rather than the full dataset.

This analytical database is needed to support adtech-specific functions such as bidding for ad slots; both A/B and multivariate testing of ads, content, positioning, and surrounding content; and quantifying and reporting ad effectiveness to advertisers. Inherent in all of this is the ability to make changes instantaneously in response to ad performance: to pull or re-content underperforming ads, or to reallocate ad expenditures to high performers.

Data speeds are highly relevant to this discussion, and Figure 3 shows the three data speeds recognized across the industry today.

¹ Massively parallel processing

Figure 3: Industry-accepted Data Speed Definitions

Speed Category	What It Means: Analytic Insights in...
Real-time Analytics	Hundreds of milliseconds (usually 100-200ms)
Near-real-time Analytics	Hundreds of milliseconds to 10 seconds
Batch Analytics	Anything longer than 10 seconds

Source: Stratecast

Given that ad decisions and placements must now be executed in real-time and near-real-time, adtech demands the ultimate in speed and performance.

Why HPE Vertica Offers a Solution Adtech Needs

In its ongoing research of the BDA market, Stratecast has identified HPE Vertica as an analytical database that delivers the speed, scale, and querying capabilities adtech requires. Figure 4 lays out characteristics of the HPE Vertica solution that make it a strong fit for adtech.

Figure 4: HPE Vertica Characteristics Essential to Adtech

Area	Where and Why Near-Real-time Insights Are Needed
Stores and manages data differently than relational DBs	<ul style="list-style-type: none"> Column store DB enables database administrator to sort, compress, and organize data on disk much more efficiently DBA also gains ability to apply specific compression and encoding algorithms to different parts of the data for efficient data sorting
Handles data queries differently	<p>Sample query, all customers who bought a certain product in a given state or nation:</p> <ul style="list-style-type: none"> Simply reads two columns off disk Creates no indexes, and self-indexes data via sorting and encoding <p>These architectural differences speed queries by orders of magnitude</p>
Provides ultra-high-speed performance	<ul style="list-style-type: none"> Runs queries, against databases with petabytes of data, 50-1,000x faster than legacy systems Scale-out MPP architecture processes queries in parallel over multiple processors Provides linear scalability and high availability
Offers Database Designer application	<ul style="list-style-type: none"> Provides instant answers to questions Automatically creates customized projections Saves time by eliminating need to create data schemas manually
Has specific built-in approximate query processing (AQP) analytical functions to support adtech	<ul style="list-style-type: none"> Supports functions including pattern matching, approximate count distinct and approximate count distinct synopsis Replaces hundreds of lines of code (and associated time/cost) of other available solutions with just five lines of code

Area	Where and Why Near-Real-time Insights Are Needed
Supports other adtech essentials	<ul style="list-style-type: none"> • 10-20-millisecond ad bidding windows • A/B and multivariate testing • Near-real-time reporting to reallocate ads and budgets on the fly

Source: Stratecast

Some of this is pretty standard, in that the same could be said not just of HPE Vertica but of any column store analytical database, as opposed to traditional relational databases. However, the product's high-performance architecture and innovative features such as its built-in AQP analytical functions make it a fit for the millisecond business that adtech has become.

Where Competitive Solutions Fall Short

Even if they have what they thought was a system that would provide useful data to inform marketing processes, ad providers are telling Stratecast many of these systems are falling short of expectations. Some solutions vendors are still offering legacy relational DB-based solutions and trying to simply throw more compute power to make them run faster. Others offer cloud-based solutions that provide flexibility and rapid rampup, but can introduce latency and simply do not have the data processing capacity to keep up with the pace of the ad business. What would take just a few lines of code in HPE Vertica takes hundreds of lines of code in these competitive systems—and HPE Vertica also offers more rapid query execution.

For these reasons, some companies in the ad ecosystem have tried other cloud-based solutions, “hit a wall,” and reverted back to the power and reliability of HPE Vertica. Since a number of adtech customers do not want to own and manage infrastructure and are comfortable with solutions in the cloud, HPE Vertica provides choice: it can be deployed as an on-premises solution, in the cloud, or on Hadoop.

HPE Vertica Covers the Bases with SQL-on-Hadoop Solution

Perhaps the most important development in the BDA space over the past 12 months, and one Stratecast has analyzed recently,² is the emergence of solutions that enable companies to apply the expressive query power of SQL, not just against traditional structured data but also now against unstructured and semi-structured data. These so-called SQL-on-Hadoop and SQL-on-Everything solutions enable users to run SQL queries against all data. Unfortunately, up to now, customers who have relied on some of the cloud-based systems in use in adtech have been faced with a poor choice:

- Deploy and manage an on-premises system to gain rapid data access, but in the process, consume precious internal resources, human and otherwise; or
- Use a cloud-based system, which conserves internal resources—but also slows DB response times and can impact their ability to conduct business.

² Stratecast, *BDA in the Enterprise: SQL Still Rules—but What Is the Impact on Privacy?* (BDA 4-10, December 2016), available [here](#)

Advertising companies with data lakes on Hadoop and hybrid SQL/NoSQL environments now have an HPE option for this: Vertica for SQL on Hadoop. Vertica for SQL on Hadoop supports open source and hybrid SQL/NoSQL environments, including other Apache technologies such as Kafka, Spark, and Storm, all of which support near-real-time analytics. With Vertica for SQL on Hadoop, an ad provider can now use SQL to query any of its DBs, and can leverage all dashboards it may have built up to now.

Core Business Benefits

Stratecast has determined that HPE Vertica offers these core business advantages and opportunities to the advertising vertical:

1. Technical and Organizational Advantages

- Features data and application architecture that supports Big Data and near-real-time analytic insights essential for digital advertising.
- Built on an open and extensible architecture to adapt to evolving marketing needs.
- Performs very fast and deep analytics on massive amounts of data associated with online advertising.
- Eliminates bottlenecks that delay getting analytic insights to users.
- Provide an easy onramp for data scientists to quickly obtain actionable information they need to do their jobs.
- Simplifies how users can now relate things that have a business relationship to each other (but may not necessarily be reflected in data model).

2. Cost-Savings Opportunities

- Helps each of the six groups outlined earlier in paper make faster, better decisions.
- Offers savings as optimized, cost-efficient ad expenditures add up fast.
- Enables closed-loop analysis to optimize marketing spend.

3. Revenue Opportunities

- Hitting optimal target audiences with right ads and messaging drives greater conversions and revenue across the board.
- Offers ability to analyze and monetize visitor behavior through A/B and multivariate testing to deliver enhanced revenue results.
- Enables more revenue generation by increasing inventory of advertising and cross-promotional offers.

4. Customer Retention Advantages

- Optimizes online experiences through personalized 1:1 user engagement.
- Predicts user engagement patterns to enable consistently high-value content.
- Enhances satisfaction, profitability, and loyalty.
- Reduces churn.

With Vertica for SQL on Hadoop, an ad provider can now use SQL to query any of its DBs, and can leverage all dashboards it may have built up to now.

Case Study: adMarketplace

Some organizations have hesitated making the badly needed upgrade to a platform that can truly handle Big Data. This is due in part to concerns that newer analytical databases, while powerful, may not be as scalable as incumbent systems; or because, plainly put: “Changing databases is hard.” Retiring an existing system and deploying a new one can seem painful. The case study and snapshots

that follow, however, provide real-world examples of technological and business benefits accruing to those in the adtech community who have stepped up and made this beneficial change.

HPE Vertica Powers Real-time Bidding Engine at adMarketplace

adMarketplace offers the largest search advertising solution of its kind after Google, Bing, and Yahoo!, and the only programmatic³ marketplace for search:

- adMarketplace's Advertiser 3D predictive ad platform provides end-to-end search advertising solution for major brands and publishers. Advertiser 3D serves up performance ads to targeted audiences. Whereas brand advertising is about boosting a company's image, performance advertising is all about creating conversions: achieving a sale or other results desired by the brand/advertiser. In short, performance ads must perform, must show immediate results.
- Advertiser 3D displays ads by using complex algorithms in concert with BDA to present an offer tailored to an individual's tastes and interests. The platform supports hundreds of advertisers through the algorithmic bidding and matching system BidSmart, which integrates with Advertiser 3D analytics to ensure that ads are matched, displayed, analyzed, and recorded in real time. BidSmart predicts relevancy, competitive bid landscape, and click value for advertisers.

HPE Vertica is the analytics engine powering Advertiser 3D because it offers the speed and performance to meet adMarketplace's stringent latency requirements for data analytics:

- Every time a spot opens for an ad on a site, a real-time auction occurs there in milliseconds based on user, location, and device type. Each request is matched to the millions of available ads and keywords, from hundreds of advertisers.
- Advertiser 3D queries execute in less than two seconds, so advertisers can access their data in real time. adMarketplace also uses data stored in Vertica for their algorithmic bidding system, BidSmart, HPE Vertica loads nearly two terabytes of data per hour into Advertiser 3D's analytics engine and responds to 500 million search queries per day

Business Results: adMarketplace and its Customers

adMarketplace achieved 118% year-over-year performance ad revenue growth in its first year after launch, and HPE Vertica has helped extend success to adMarketplace's own customers, who rely on the Advertiser 3D platform:

Mazda – adMarketplace has increased leads for Mazda by 38% at a lower cost than Google's Search Partner Network.

Lenovo - adMarketplace is driving more than half the volume of search-driven traffic to the manufacturer as search behemoths Google and Bing-Yahoo!—combined.

Netflix – adMarketplace generated month-over-month mobile and tablet conversion rate increases of 36%.

³ Fully-automated

Case Study Snapshots

Figure 5 presents case study snapshots reflecting either technical or qualitative business enhancements adtech clients are experiencing as a result of deploying HPE Vertica.

Figure 5: Clients’ Technical and Qualitative Business Benefits from Deploying HPE Vertica

Client	Description/Impacts/Results
AdForm	<ul style="list-style-type: none"> • More than 21,000 advertisers, 1,600 agencies and 800 of the world’s leading publishers currently employ Adform technologies and services; company is growing 174% YoY • HPE Vertica and Adform platform handle 1+ million requests per second; each requires complex statistical/modelling/decision-making calculations on many datasets including each user’s digital journey touchpoints, CRM data, and third-party demographic profiles • With HPE Vertica, clients have the flexibility to perform <i>ad hoc</i> queries in approximately two seconds and 99% of all reports are generated within three seconds. These near-real-time insights help AdForm meet customer service-level agreements (SLAs)
AdGear Technologies, Inc.	<ul style="list-style-type: none"> • Full-stack ad platform offers advanced advertising analytics, proprietary targeting, cross-channel attribution measurement, ad serving, real-time bidding (RTB) and exchange technology • Implementing customer analytics and campaign management with HPE Vertica; currently managing 1-5 terabytes of data and using both Hadoop and homegrown business intelligence (BI), visualization, and ETL tools • With HPE Vertica, AdGear is achieving faster analysis and reporting; a more interactive, conversational view of its data; the ability to analyze more, and more complex, data; the ability to readily add additional database capacity; and high availability • AdGear has increased query performance response time 5-10x with HPE Vertica
AppNexus	<ul style="list-style-type: none"> • RTB exchange based on HPE Vertica for advertisers and online content providers • Real-time ad serving platform, advanced yield management controls, optimization algorithms and patented brand and safety monitoring • Supports hundreds of companies worldwide buying and selling billions of online ads • Criteo and Simpli.fi—also HPE Vertica customers--use AppNexus and other similar exchanges for bidding purposes, and use HPE Vertica for more comprehensive campaign reports for brands/advertisers
Cardlytics	<ul style="list-style-type: none"> • Proprietary native online and mobile banking channel enables advertisers to deliver relevant offers to financial institution customers based on actual purchase behavior • With HPE Vertica, Cardlytics is achieving a 10x increase in average customer pipeline due to 40-80x faster query performance; and a 90% reduction in operational support, with four full-time employees repurposed to value-added analytics activities • Cardlytics realized three-month ROI with HPE Vertica
DoubleVerify	<ul style="list-style-type: none"> • Provides media authentication solution to guarantee authentic impressions and eliminate ad fraud • HPE Vertica’s data compression and high performance is enabling DoubleVerify to handle huge increases in traffic and fraud analysis demands
SimpleReach	<ul style="list-style-type: none"> • HPE Vertica-based system provides near-real-time visibility and detailed historical reporting of content performance • Metrics include reach, engagement, and social activity

Client	Description/Impacts/Results
Tapjoy	<ul style="list-style-type: none"> • HPE Vertica-based system handles 1.5 million interactions per day • Supports 200 simultaneous system users viewing analytic insights • Supports Microstrategy, Tableau, and <i>ad hoc</i> queries
(NDA) Provider of global technology platform for ad buyers	<ul style="list-style-type: none"> • Assists buyers in RTB space for advertising exchanges, managing ad placement on high-traffic customer Web sites; primarily stores RTB and non-bidding data in Hadoop • Using HPE Vertica's FlexTables for analytics against "hot data" (RTB data), and as the analytics engine to explore non-RTB data for <i>ad hoc</i> queries and more
(NDA) Provider of unified campaign reporting solutions	<ul style="list-style-type: none"> • Generates real-time ad analytics and content analytics • Integrates its own ad and content analytics with inbound streaming data from DoubleClick, Sizmek, Atlas, OAS, AppNexus, PointRoll, Nielsen OCR/Digital Ad Ratings, BlueKai, Krux, Lotame, and Grapeshot • Provider's HPE Vertica-based platform combines and analyzes all internal and external data to deliver campaign reporting and single-click PowerPoint automated snapshots for brands/advertisers

Sources: HPE Vertica and the companies

Stratecast The Last Word

The \$600 billion global advertising market has long benefited from John Wanamaker's famously humorous statement that he knew half his advertising spend was wasted, but could never figure out which half. **In other words, the advertising industry benefited from a considerable amount of waste spending**, but advertisers were willing to put up with it because results were sufficient to justify continuing such practices—and because they were limited in their ability to pin down which ads were delivering the best results.

Those days are long gone. Today, sophisticated ad attribution technologies, not to mention a sophisticated new breed of advertisers, are forcing providers to quantify, in real time, which ads are generating conversions and which are not—and demanding that they revise ads, advertising venues, and placement among surrounding content on the fly to improve results. Making this quantum leap a reality means supporting six key constituent groups that drive digital ad processes, and handling hundreds of millions of queries and other interactions per day. This was well beyond the capabilities of the existing platforms most ad networks and providers were using, and even modern cloud-based and open source engines have been found unworthy of the task.

HPE Vertica is an analytical database that delivers the speed, scale, and querying capabilities required to keep up with the frenetic pace of the world where digital ads are bid, bought, sold, and placed. HPE Vertica manages and analyzes massive amounts of data to provide the rapid analytic insights that both humans and machines need to exist and compete on this ultra-high-speed playing field. As HPE continues to prove itself in the digital realm with providers and brands including AdGear, adMarketplace, AppNexus, and TapJoy, it will continue to create new opportunities for itself. Perhaps the best news of all is the upside: as hard as digital is surging, digital ad spending still only accounts for about a third of total ad spend.⁴ That leaves plenty of fertile digital ground to cultivate for the next decade or so.

Jeff Cotrupe

Industry Director – Big Data & Analytics

Stratecast | Frost & Sullivan

jeff.cotrupe@frost.com

⁴ eMarketer, *US Digital Ad Spending to Surpass TV this Year*, as cited [here](#) and available [here](#)

About Stratecast

Stratecast collaborates with our clients to reach smart business decisions in the rapidly evolving and hyper-competitive Information and Communications Technology markets. Leveraging a mix of action-oriented subscription research and customized consulting engagements, Stratecast delivers knowledge and perspective that is only attainable through years of real-world experience in an industry where customers are collaborators; today's partners are tomorrow's competitors; and agility and innovation are essential elements for success. Contact your Stratecast Account Executive to engage our experience to assist you in attaining your growth objectives.

About Frost & Sullivan

Frost & Sullivan, the Growth Partnership Company, works in collaboration with clients to leverage visionary innovation that addresses the global challenges and related growth opportunities that will make or break today's market participants. For more than 50 years, we have been developing growth strategies for the Global 1000, emerging businesses, the public sector and the investment community. Is your organization prepared for the next profound wave of industry convergence, disruptive technologies, increasing competitive intensity, Mega Trends, breakthrough best practices, changing customer dynamics and emerging economies? For more information about Frost & Sullivan's Growth Partnership Services, visit <http://www.frost.com>.

CONTACT US

For more information, visit www.stratecast.com, dial 877-463-7678, or email inquiries@stratecast.com.