

# ThingSpan™: The Distributed Graph Platform for Digital Marketing

## The Business Opportunity

The digital marketing industry is rapidly evolving, both in terms of its business requirements and the enabling technologies needed to improve decision-making and gain competitive advantage. While the ad tech community has effectively utilized NoSQL technology, its concentration has been on meeting the speed requirements with aggregate-oriented key value stores and column family databases. To gain and maintain a competitive advantage in this fast-moving industry, it is crucial for digital marketing to adopt a massively scalable distributed graph platform that can handle:

- Graph analytics and real-time relationship discovery
- Integration with the open source stack – Spark, Kafka, HDFS, YARN
- High-speed ingest with parallel querying
- Petabyte-scale to trillions of nodes and edges

“ Forrester Research estimates that graph databases will be the fastest-growing area in database management systems, with more than 25% of enterprises using graph by 2017<sup>1</sup>. ”

An enterprise-grade graph analytics platform can be used to augment these three broad categories:

### · **Lookalike Audience Targeting**

Big data analytics is crucial for targeting the right audience and using all available information to widen customer reach. A scaleable graph analytics platform lets digital marketers identify “lookalike audiences” or groups of potential customers who share certain parameters with the ideal target customer. First, second, and third-party data sources are combined in near real time to create wider segments prioritized by value to efficiently scale marketing campaign reach to larger populations.

### · **Cross-device marketing**

Big data analytics is crucial for evaluating impression opportunities and brand engagement as the number of internet-connected devices is exploding. A graph analytics platform lets digital marketers expand customer reach and more accurately measure advertising effectiveness by seamlessly connecting the dots between devices to gain a completely unified view of their target audiences.

### · **Advanced attribution modeling**

With so many ways to reach customers, it can be difficult to attribute campaign success to the right channel. Statistical averages now fall short, so using a graph analytics platform to uncover insights deep within your marketing data from all available data sources is the key to continually optimizing digital spend and increasing ROI.



## Use Case: Lookalike Audience Targeting

Through sub-graph similarity algorithms, digital marketers can identify lookalike audiences segmented by the percentage of matching criteria and use this information to tier their spend appropriately across the digital ecosystem.

In addition to lookalike audience targeting, Objectivity’s ThingSpan scaleable graph analytics platform is ideally suited for device matching, dwell detection, and content recommendation use cases. Best of all, ThingSpan integrates with your existing NoSQL or NewSQL technology stack.

<sup>1</sup> Big Data Analytics News:

<http://bigdataanalyticsnews.com/4-predictions-for-nosql-technologies-in-2016/>

## The Technical Challenge

Many organizations rely on big data analytics solutions that involve a data ingestion layer that processes all data points about customer data—including ecommerce clickstreams, point-of-sale data, and social media—before breaking them down into stored memory where they can be queried. However, while large volumes of data may be processed in this manner, businesses often rely on micro-batches and may need to wait days until the relevant data points can be surfaced. In many cases—particularly when dealing with profit opportunities and customer churn—such delays are not an option.

In order to generate the high-performance, high-speed processing power and the sophisticated contextual analysis needed within the digital marketing industry, an enterprise-class graph analytics solution is essential.

The challenge, however, is finding a distributed graph platform that can scale to petabytes of data and perform queries in parallel to data ingestion. While there are many platforms available, few of them offer real-time data analysis at scale, enabling organizations to capture streaming data and analyze it in relation to historical and contextual data to gain a 360-degree view of their customers across a broad array of use cases.

Digital marketers need a highly scalable, real-time graph analytics platform to analyze massive volumes of advertising and customer data for campaign effectiveness, business performance, predictive analytics, and other benchmarks that are high priority to organizations.

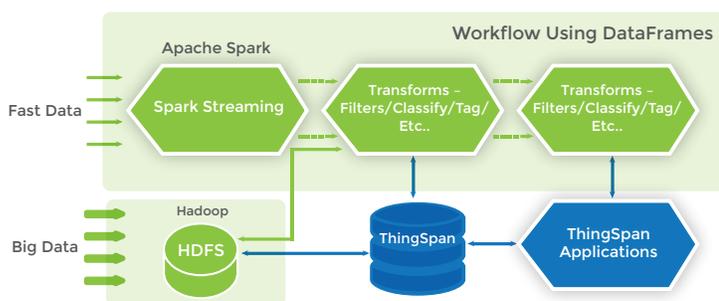
## The ThingSpan™ Solution

ThingSpan™, Objectivity's massively scalable graph analytics platform, is compatible with the distributed open source data management framework, **Hadoop**, and utilizes **Spark** for high-speed streaming ingest and enriched data processing. It has the power to transform and analyze real-time advertising and customer data in context, offering a single logical view of all data—wherever located—to accelerate parallel processing. ThingSpan's subgraph similarity capability helps find emerging patterns with high degrees of accuracy.

ThingSpan leverages open source tools by supporting the Hadoop and Spark ecosystem atop a high-performance, distributed graph platform purpose-built for relationship and pattern discovery. It runs natively on top of POSIX or HDFS as a YARN application while using Spark for workflow and data transformation. It is also designed to support streaming systems based on Kafka, Flume, and other distributed messaging tools for streaming data. Integration with other NoSQL solutions and with Spark via DataFrames allows ThingSpan to ingest streaming data while maintaining and persisting relationships as first-class logical models.

This model allows for enriched and transformed data to simplify the support of complex, multi-dimensional queries associated with digital marketing applications and analytics. ThingSpan enables businesses to capture powerful insights around data relationships to make better-informed decisions in media buying and campaign management, avoid customer churn, and discover new ways to reach audiences across all channels and devices.

### ARCHITECTURE DIAGRAM



Objectivity, Inc. delivers massively scalable and highly performant distributed database platforms that are proven to power mission-critical applications for the most demanding and complex datasets in the enterprise. Objectivity helps organizations to rapidly build new Spark Streaming-enabled solutions for finding connections and patterns through graph analytics within petabytes of data, stored in HDFS, to achieve real-time relationship discovery. With a rich history serving Global 1000 customer and partners, Objectivity holds deep domain expertise in fusing vital information from massive data volumes to capture new revenue opportunities, drive competitive advantages and deliver better business value. Objectivity is privately held with headquarters in San Jose, California. Visit <http://www.objectivity.com> to learn more.



ThingSpan is a massively scalable distributed graph platform designed specifically for the complex issue of extracting actionable insights from Fast Data and Big Data sources to enable real-time relationship discovery. It is architected to integrate with major open source Big Data technologies, such as HDFS and Apache Spark. ThingSpan leverages Objectivity's core object and graph data-modeling technologies and the company's rich experience in building fusion solutions.