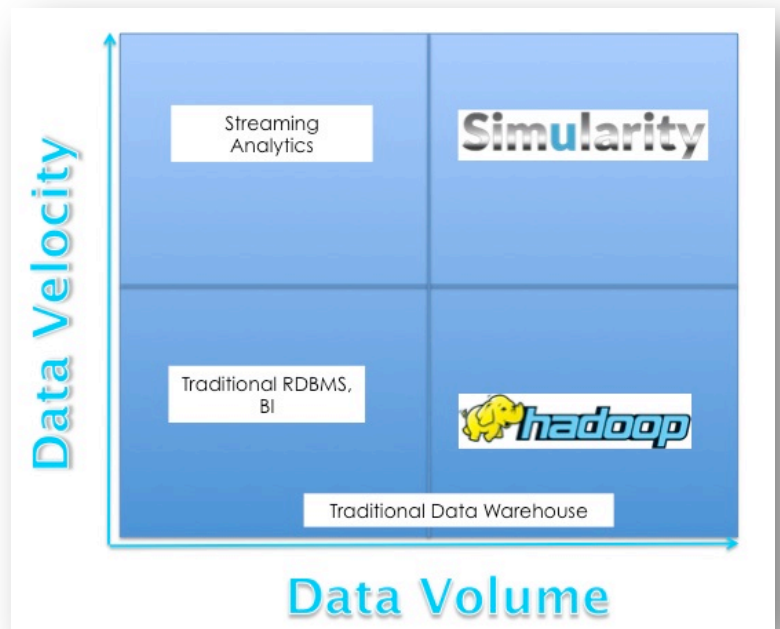


## Turbo Charge Your Analytics With Similarity

Similarity provides the first and only real time similarity and correlation analytics software for truly huge datasets.

- o **Integrate** with ease using our web services API
- o **Delight** your users with real time responsiveness
- o **Detect** hidden patterns of correlation faster than ever
- o **Predict** churn, failure, or fraud in real time with our dynamic classifier
- o **Replace** your home grown analytics with our powerful and simple solution



## The Right Tool For The Job

Similarity's High Performance Correlation Engine (HPCE) is purpose-built for correlation and similarity analytics on truly huge data sets. It's not a general purpose data management solution or statistical package. Our graph-like in-memory data representation is extremely compact and optimized for correlations. Our massively parallel processing strategy is more powerful and efficient than Map-Reduce. Our Dynamic Classifier leverages the HPCE to create archetypes used for predictive similarity analysis.

## Size Matters

**Bigger is Better** The more data you can include in your analysis, the stronger your conclusions will be. When you've got sparse data, or data from a wide variety of sources, the true benefits of analytics are only realized when you can examine all of the data at once. Got hundreds of millions of rows in your data set? No problem. Our response times don't go up when you add more data, and our ability to scale is truly linear.

Our response times don't go up when you add more data, and our ability to scale is truly linear.

"Similarity gives us an edge that other recommendation services just can't offer: better and faster recommendations, powered by all our data, on a single inexpensive server."

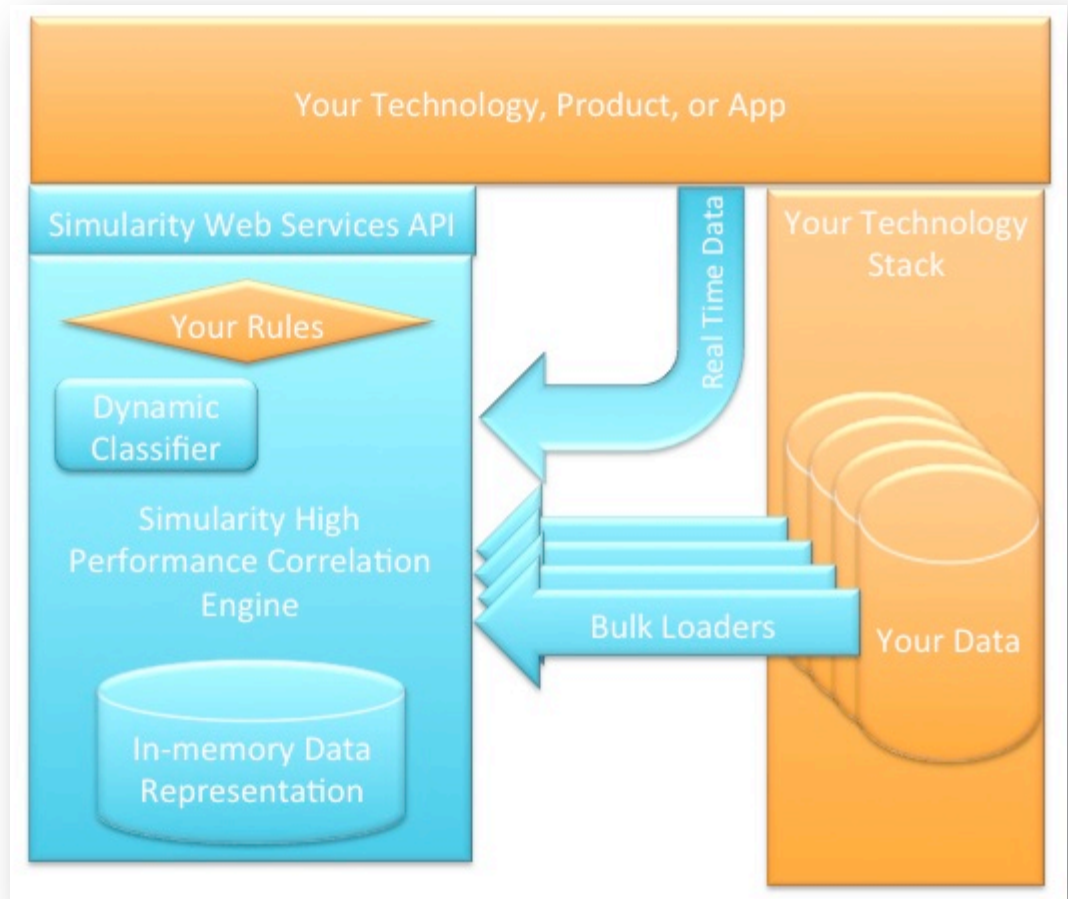


--Casey Carey  
VP of Marketing  
Alibris

**Smaller is Smarter** Big data doesn't have to mean big expenses. When it comes to cost, Similarity helps you keep it small. Our amazing efficiency allows you to do blazingly fast analytics with 90% less hardware. That means fewer servers, but it also means smaller monthly bills for power consumption or hosting, and less IT overhead. You can even reuse the commodity servers you already have – Similarity runs on any Linux server.

## Your Data, Your Way

Similarity's data representation is super flexible (it's a modified triple store), so it's easy to load your data, no matter where it lives or what form it takes. You can create custom rules that can apply to everything from data ingestion to result filters. You can use our standard loaders or write your own in any language. Load your data via the API or over a socket. Add or remove data sources at any time. We give you the flexibility to do it your way.



## Get Real, Fast

Real-time data ingestion and analysis is built-in. Our compact internal data representation means no time consuming indexes to build and no pre-calculations. You get high velocity data analysis with full historical context. When you incorporate Similarity as your analytics layer, your product or app can do real-time, faceted, ad hoc analytics over trillions of data points with response times measured in milliseconds.

## Leave The Analytics To Us

Finally there is an easy answer to the buy vs. build question for your big data analytics. Now it's easy to replace your homegrown solution with an affordable out-of-the-box solution that gives you unparalleled performance on huge datasets. Integration is easy and the Total Cost of Ownership can't be beat. The choice is clear. Let us do the heavy lifting, so you can concentrate on building your business.



[www.similarity.com](http://www.similarity.com)

*Embedded High Performance Analytics*

[info@similarity.com](mailto:info@similarity.com)

[@similarity](https://twitter.com/similarity)