



Akamai Improves Ad Targeting through Faster Analytics

Soaring Data Volumes and a High Computational Load

Akamai is a community of retailers, brand marketers, and consumers. It helps advertisers deliver more targeted, more relevant messages based on actual consumer shopping behavior.

Because historical shopping data is the most reliable predictor of future behavior, Akamai aggregates this data, applying sophisticated predictive modeling analytics to give web advertisers improved advertising performance, and to provide web publishers with higher yielding placements.

As this sophisticated network has grown to encompass more than 450 brands and multi-channel Internet retailers, delivering the right ad to the right audience at the right time has become far more challenging—with data processing bogged down by soaring data volumes and high computational loads.

Speed Up Processing of up to 75 Million Daily Events

Akamai's core business value relies on mining a huge set of data that represents up to 75 million daily events using models that provide advertisers improved targeting to boost revenue and yield.

These models are extremely sensitive to tight time windows because consumers tend to act quickly once they've identified themselves through their online activity as prospects for specific products and services.

With more than 140 million anonymous profiles, and 230 models run for each record, Akamai faced constant challenges scoring all these records—daily reports on analysis and scoring were delayed by up to 20 hours.

"With shopping behavior of over 140 million consumers to analyze daily, increasing data volumes were pushing our analysis latency to unacceptable levels," said Peter Kools, chief architect at Akamai.

The goal? Reduce this 20-hour lag time to six hours or less to give advertisers and media partners quick access to vital information.

As Andrew Nantkes, data warehouse architect at Akamai, noted, "Anything Aster Data can do to cut analytics time for us is a major business win."

Akamai Background

- Industry's only online shopping data co-op
- 450 brand and multi-channel internet retailers
- Data on 140 million anonymous U.S. online shoppers

Challenges

- Boost ad revenue and yield for retailers
- Quickly score anonymous individuals' daily activity
- Refresh models daily, analyzing up to 100+ million users against 1000's of attributes over 100's of behaviors

"With Aster Data nCluster, we've been able to cut our analysis turn-around time by 70%, greatly increasing the relevance and value of our targeted advertising."

*Peter Kools, Chief Architect
Akamai*

Always-On, Always-Parallel MPP Architecture

Aster Data *nCluster* is the industry's first massively parallel (MPP) data warehouse architecture that allows applications to be fully embedded within the database engine to enable ultra-fast, deep analysis of massive data sets.

A significant improvement over traditional data warehouses that were never designed to keep up with the data loads of today's big data applications, Aster Data's massively parallel data-application server effectively uses Aster Data's patent-pending SQL-MapReduce with data and application parallelization to address the big data challenge.

Thanks to an always-on, always-parallel MPP architecture, Aster Data *nCluster* overcomes the challenges of slow data load and excessive query times through techniques that avoid network bottlenecks inherent in traditional data warehouses.

Aster Data deployed its *nCluster* data-application server across Akamai's Dell commodity hardware cluster with multiple terabytes of capacity.

"Blazing Fast" Performance, Improved Analytics

After deploying the *nCluster* data-application server, members of the Akamai ad network were able to take quick advantage of improved analytics to increase ad revenue and yield.

Almost immediately, the time required for analytic record generation dropped from 18 to 20 hours to less than six hours—with new business-level reports run in just 38 seconds—a level Akamai executives had once deemed "inconceivable."

"With Aster Data *nCluster*, we've been able to cut our analysis turn-around time by 70%, greatly increasing the relevance and value of our targeted advertising," explains Peter Kools, Akamai's chief architect. "Aster Data has given us the ability to do analytics we couldn't do before, and Aster Data's ability to count distinct things over time periods is blazingly fast."

And as the company's director of network operations, Wayne Earl, points out, "It really does what you said it would do. I've heard just about every story in the industry from vendors on what they could do. I was a cynic on what Aster Data was claiming to do. The fact that Aster Data not only met, but exceeded my expectations was unbelievable."

About Aster Data

Aster Data is a proven leader in big data management and big data analysis for data-driven applications. Aster Data's *nCluster* is the first MPP data warehouse architecture that allows applications to be fully embedded within the database engine to enable ultra-fast, deep analysis of massive data sets. Aster Data's unique "applications-within™" approach allows application logic to exist and execute with the data itself. Termed a "Data-Analytics Server," Aster Data's solution effectively utilizes Aster Data's patent-pending SQL-MapReduce together with parallelized data processing and applications to address the big data challenge. Companies using Aster Data include Coremetrics, MySpace, comScore, Akamai, Full Tilt Poker, and ShareThis. Aster Data is headquartered in San Carlos, California and is backed by Sequoia Capital, JAFCO Ventures, IVP, and Cambrian Ventures, as well as industry visionaries including David Cheriton, Ron Conway, and Rajeev Motwani. For more information please visit www.asterdata.com, or call 1.888.Aster.Data.

Solution

- Aster Data *nCluster* massively parallel data-application server
- Dell PowerEdge 2950 servers
- Replication factor: 2

Results

- Faster analytic record generation (from 20 hours down to six hours)
- Complex business-level reports for clients run in just 38 seconds
- Reduced execution of key data mining process by 70% for improved analytics

"[Aster Data's] ability to count distinct things over time periods is blazingly fast."

In addition, Aster Data's team and support have been outstanding."

*Peter Kools, Chief Architect
Akamai*