



fueled by Aerospike

## Snapdeal, India's Largest Online Marketplace, Accelerates Personalized Web Experience With Aerospike

In India, Internet retail accounted for \$3 billion in 2013 and is expected to swell to \$22 billion in the next five years. Still, with a total consumer market valued at \$450 billion, India has been called the world's last major frontier for e-commerce. Leading the industry in addressing this opportunity is Snapdeal, India's largest online marketplace.

Snapdeal empowers sellers across the country to provide consumers with a fully responsive and intuitive online shopping experience through an advanced platform that merges logistics subsystems with cutting-edge online and mobile payment models. The platform has a wide range of products from thousands of national, international and regional brands.

Snapdeal.com now has a network of more than 20,000 sellers, serving 20 million-plus members—one out of every six Internet users in India. Powering this platform is the Aerospike flash and DRAM-optimized in-memory NoSQL database. By harnessing the real-time big data processing capabilities of Aerospike, Snapdeal provides consumers with the up-to-the-moment updates about product availability, pricing, and seller ratings to help make buying decisions informed and convenient.

**"There has been no need for maintenance with Aerospike database; it just works out of the box."**

**Amitabh Misra**  
Vice President of Engineering,  
Snapdeal

### The Demands of Dynamic Updates

Snapdeal's business and platform model is anchored by an innovative system that incorporates updates from both sellers and shoppers. The Snapdeal.com platform enables sellers to list products for sale on the site, manage inventory, and make pricing changes in real-time based on what is happening in the marketplace. High volume—for example, a pair of shoes sells every 30 seconds—means that thousands of sellers are making dynamic price adjustments. This results in Snapdeal's inventory and pricing management system processing more than 500 writes per second.

Meanwhile, shoppers on Snapdeal.com review and rate sellers for customer satisfaction based on their experience with the product, including shipping, delivery and returns.

## ◀ AEROSPIKE ▶

### In-memory NoSQL Database

#### RESULTS

- In-memory Aerospike database maintains sub-millisecond latency on Amazon Elastic Compute Cloud (EC2) while managing 100 million-plus objects stored in 8GB of DRAM to support real-time dynamic pricing.
- Predictable low latency with 95-99% of transactions completing within 10 milliseconds—essential for enabling a responsive customer experience.
- Aerospike's highly efficient use of resources enables Snapdeal to cost effectively deploy on just two servers in Amazon EC2.
- Full replication across the Amazon EC2 servers ensures business continuity.
- Aerospike's ability to work out of the box enabled Snapdeal to use Aerospike Client APIs and get up and running in days and nearly eliminate the need for maintenance.

#### COMPANY PROFILE

**Website:** [www.snapdeal.com](http://www.snapdeal.com)

**Industry:** Retail, Internet Marketplace

**HQ:** New Delhi, India

#### GOAL

Deliver real-time dynamic pricing while maintaining consistently high availability and responsiveness for millions of consumers purchasing products from 20000+ sellers.

With every page click, Snapdeal.com combines the updates from shoppers and sellers to display the most relevant products, as well as rankings for all the sellers that are offering the product by price, delivery time, and customer satisfaction.

The challenge for Snapdeal has been maintaining the platform's real-time performance even as the company's growth accelerates.

"We have built up both the business and the technology platform in just two years," said Mr. Amitabh Misra, Snapdeal vice president of engineering. "We make an average of two releases a day and sometimes we turn up seven or eight releases in a day. In two years, we have scaled more than 200 times—the number of products listed, the number of sellers we have, the amount of business they do, the number of servers, storage and the technology team—everything has grown 200 times."

From a seller standpoint, Mr. Misra explained, "An efficient marketplace requires that sellers to be able to push their updates in real-time. As more sellers sign up and more products are listed and more price changes are made, we knew we needed to scale the system."

At the same time, Snapdeal has had to address the consumer demands on the Internet marketplace.

"Studies have shown that slow-loading pages traditionally result in shopper abandonment," said Mr. Amitabh Misra, Snapdeal vice president of engineering. "We determined that in order to support the business, we needed our system response times to complete within 10 milliseconds, 95% to 99% of the time—even while our platform processed up to 10,000 reads per second."

### Success Leads to Growing Pains

To support its inventory and pricing system, Snapdeal initially deployed 10 MongoDB NoSQL database servers with 5 GB of data in DRAM as a cache in front of MySQL. The Snapdeal application used write-through techniques to update information first in MySQL and then in MongoDB, and it processed reads from MongoDB.

However, as the business scaled and more sellers made price adjustments on more products, the MongoDB response times shot up from 5 milliseconds to more than a full second. This not only compromised the consumers' shopping experience; it also led to lost revenue opportunities.



"Mongo was working perfectly fine. It was giving us very consistent response times, and it scaled well with hardware," Mr. Misra recalled. "But as our seller base swelled and they started listing more products and making more changes in response to the dynamics of the marketplace, we had to moderate the rate at which we were pushing updates. We had to spread them out evenly throughout the day as opposed to being able to absorb them in near real-time. Sometimes the updates had to wait for hours," Mr. Misra recalled.

He added, "We ran tests to make sure that prices could be re-lected in near real-time when millions of buyers a day were loading pages. We were getting stuck when there was a surge in concurrent price updates from many sellers and we saw degradation in the buyer experience."

With 500% growth in 2013 and revenue projected to exceed \$1 billion by 2015, Snapdeal realized that it would require a high-throughput, low latency system that could scale from 2 million to 150 million records, from 10,000 to 30,000 reads per second, and from 30 to 500 writes per second.

**"With Aerospike, we can push through huge price changes while maintaining the same response time experience on the buyer's side—even with millions of buyers. That has been the biggest advantage."**

**Amitabh Misra**  
Vice President of Engineering,  
Snapdeal

Mr. Misra explained, "We did not want a system that gave us average performance that was very good but then degraded for 10% of the calls. That would not work for us for a real-time system."

At the same time Snapdeal wanted a technology solution that could affordably scale as the business expanded. Notably, operational efficiencies directly impact profits, so Snapdeal.com did not want to use an expensive clustered relational database management system (RDBMS).

"We decided to evaluate a variety SQL and NoSQL technologies, including Aerospike's in-memory NoSQL database," Mr. Misra said.

## Evaluating Database Options

Snapdeal looked at the option of upgrading its MongoDB NoSQL implementation. Additionally the company evaluated Couchbase, Redis, Terracotta BigMemory Max, Amazon Memcache, Amazon DynamoDB, and the Aerospike in-memory NoSQL database.

Mr. Misra noted, "We were looking for a solution that gave us three things; One: cost-effectiveness. Two: concurrent reads and writes, which can scale up with added hardware. Three: performance that was very good 95% to 99% of the time."

Most of the systems reviewed by Snapdeal engineers failed to meet at least one of the criteria:

- The existing MongoDB solution lacked predictable response times under high write loads, sharding was complex, and hardware requirements for scaling were cost prohibitive.
- Clustered RDBMS databases and the Terracotta BigMemory Max caching technology were too expensive.
- Redis did not have server side distribution mechanism at that point of time.
- Amazon Memcache and Amazon DynamoDB did not deliver predictable low latency.
- Couchbase replicas could only be used as backup copies, not to distribute load. Moreover, it required twice the number of servers as Aerospike for the same throughput.

By contrast, the Aerospike in-memory NoSQL database provided several advantages. First, it performed with predictable low latency with 95-99% of transactions completing within 10 milliseconds—essential for enabling a responsive

customer experience. Second, Aerospike had the highest throughput. Third, Aerospike delivered the highest price/performance—offering the lowest cost solution in terms of both hardware requirements and ease of operations.

“None of the other systems that we evaluated prior to Aerospike were fully viable options for us,” Mr. Misra recalled. “Ultimately, we selected the Aerospike in-memory NoSQL database because it delivered the highest performance and enabled the lowest cost of deployment. It’s a fantastic product.”

### Predictable Sub-Millisecond Performance for Dynamic Pricing

Today, Snapdeal’s advanced marketplace platform, which is written in Java, includes sub-systems for order and catalog management, inventory and pricing management, fulfillment centre management, shipping, delivery and tracking management and TrustPay, a buyer-seller protection platform.

The Java-based Snapdeal inventory and pricing management system uses Aerospike to provide predictable sub-millisecond responses while managing 100 million-plus objects stored in 32 GB of DRAM. The data stored includes seller and product IDs, inventory, seller rankings and pricing attributes. Product and price changes are made to both Aerospike and MySQL while seller rankings and product details are read from Aerospike.

**“Aerospike has really come out with flying colors, and between 90% to 99% of the time, we have been getting the same consistent numbers.”**

**Amitabh Misra**  
Vice President of Engineering,  
Snapdeal

The implementation runs on two Linux servers on the Amazon Elastic Compute Cloud (EC2), and it takes advantage of Amazon Elastic Block Store (EBS) for persistent block-level cloud storage.

“As a start-up, we saw significant advantages to implementing our marketplace in the cloud and using open source software throughout that platform,” Mr. Misra notes.

### Enabling a Predictable Buyer Experience

Since implementing Aerospike, Snapdeal has seen gains in efficiency within the organization.

“We were up and running with Aerospike in a matter of days. It was extremely easy,” Mr. Misra recalled. “There has been no need for maintenance with the Aerospike database; it just works out of the box.”

More important has been the ability to scale to meet the company’s rapidly growing business volume while ensuring a first-class consumer experience.

“With our past database, whenever there was a search in concurrent price updates from many services, we saw degradation in the buyer experience,” Mr. Misra observed. “Now with Aerospike, we can push through huge price changes while maintaining the same response time experience on the buyer’s side—even with millions of buyers. That has been the biggest advantage.”

**Read more customer success stories @ [www.aerospike.com/customers](http://www.aerospike.com/customers)**

2525 E. Charleston Rd #201 | Mountain View | CA | 94043  
+1 408.462.AERO (2376) [www.aerospike.com](http://www.aerospike.com) | [sales@aerospike.com](mailto:sales@aerospike.com)

