



Wide, Fast SATA: Scaling Up High Performance Multi-Terabyte Pure Flash Storage with the Aerospike* In-Memory NoSQL* Database

Aerospike - the *Only* In-Memory (DRAM + Flash) NoSQL Operational DBMS.

◀ AEROSPIKE ▶



appnexus

"AppNexus* operates at massive scale while paying close attention to the economics of the platform. Aerospike's flash optimizations running on top of Intel® SSDs have given us the price, performance, reliability, and serviceability we need to grow our business."

- Timothy G Smith
SVP Technical Operations,
AppNexus

The Aerospike distributed database uses DRAM and Solid-State Drives (SSDs) or Flash PCIe* cards for predictable high performance, mission-critical applications. Aerospike scales up by managing indexes in DRAM and data in DRAM and Flash, and scales out to consistently process over 200k transactions per second per node with sub-millisecond latency. Aerospike is the first database to demonstrate how Flash and SSD technology can be used as part of an in-memory computing architecture for real-time big data driven applications.¹

Aerospike's customers typically require high transactional performance with high availability for multi-terabyte data sets. Aerospike's SSD optimized database has seen the progression from 100 GB SATA drives (4 per server) to larger SATA drives (12 to 20 per server node) and smaller server clusters. It enables customers with even the most demanding deployments to keep their nodes per cluster count down to a smaller, more manageable size.

CHALLENGES

- **Scaling to meet growing capacity needs.**
AppNexus' infrastructure was reaching capacity limits and a new storage solution was needed to support company growth and provide high quality service to customers.
- **Growing infrastructure costs and complexity.**
The existing infrastructure had grown to about 50 machines per cluster, resulting in growing maintenance, energy, and operational costs.

SOLUTION

- **Aerospike NoSQL Database with the Intel® SSD DC S3700 Series.**
AppNexus evaluated many SSDs and ultimately choose the Intel® SSD DC S3700 Series to run alongside its Aerospike NoSQL Database.

IMPACT

- **Greater capacity with remarkable performance, reliability and serviceability.**
AppNexus has the capacity to support business growth while providing customers with a high quality of service.
- **Reduction in cluster size.**
By integrating Intel® SSDs in a wide SATA configuration, AppNexus reduced its cluster size from about 50 machines to 8 machines per cluster, lowering maintenance, energy, and operational costs.

Spotlight on Aerospike

Aerospike offers the only flash-optimized in-memory NoSQL database that delivers predictable high performance for mission-critical, high scale applications. Customers growing their business with Aerospike include AppNexus, BlueKai*, eXelate*, The Trade Desk*, Chango* and [x+1]*. For more information, visit aerospike.com.

Spotlight on AppNexus

AppNexus, which offers the most powerful, open and customizable advertising technology platform, serves the largest and most innovative buyers and sellers of online advertising, including Microsoft* Advertising Exchange, Interactive Media* (Deutsche Telekom*), and Collective*. AppNexus is headquartered in New York City with nine global offices that employ hundreds of the brightest minds in advertising and technology. For more information, visit appnexus.com.

Learn more about how to configure a wide SATA. Visit aerospike.com/intel

Find the Intel Solid-State Drive that's right for you. Visit intel.com/go/ssd for product information.

To learn more about other Intel business solutions, please visit the Reference Room at intel.com/references.

Solution Provided By:



¹ Aerospike. www.aerospike.com/products/

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

© 2013 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others. 12/2013/FO/PAS 329854-001US

One Aerospike customer – AppNexus* – offers one of the industry's most advanced technology platforms that empowers companies to build, manage and optimize their entire online advertising business. AppNexus needed a cost-effective storage solution with predictable low latency and high throughput that could scale rapidly as business grew. AppNexus selected the Aerospike NoSQL Database and the deployment has grown to about 50 machines per cluster.

As the company looked to add more capacity and deploy the next generation of their storage system, AppNexus wanted to reduce cluster size and simplify operations. An architecture consisting of a small number of large capacity PCIe cards did not meet their targets for price, performance, reliability and serviceability. AppNexus found a SATA-based SSD solution with a large number of smaller devices delivered better performance and was less expensive. The server front-mounted drives were also easier to service compared to PCIe cards inside the server case. In addition, they were able to reduce their Aerospike cluster size from 50 to just 8 machines.

What criteria did Aerospike and AppNexus use to choose an SSD? The application and database requirements demanded consistent quality of service for the high, mixed IOPS workload. Aerospike's quality of service criteria required support for a targeted level of IOPS, while maintaining a latency threshold, which is the percentage of IOPS completed less than 1, 2 or 4 milliseconds. Additional criteria included the reliability of the drive and serviceability by their desired vendor, Dell*.

After testing many SSD's, AppNexus chose the Intel® Solid-State Drive DC S3700 Series and the Intel® Solid-State Drive DC S3500 Series in the Dell* R720xd* chassis and H710p* host controllers. This hardware configuration provided the foundation for the next generation storage system and the Aerospike deployment at AppNexus, used for managing key user information in the online advertising and real-time bidding market segments.

"Aerospike's game-changing use of DRAM and Flash showcases the potential of database applications providing microsecond user space performance with amazing consistency. Intel® SSD technology helps companies like Aerospike disrupt the database market segment, grow commercial adoption and generate business value for real-time web-scale applications."

- Rob Crooke,

Vice President and General Manager of the
Non-Volatile Memory Solutions Group, Intel

Cost-Effective Scaling with Aerospike and Wide SATA Configuration

In a wide configuration with up to 20 SSDs per 2U server, individual servers can cost-effectively serve a 16 TB/second total cluster throughput rate at roughly \$2 per GB with the Intel® SSD DC S3700 Series or roughly \$1 per GB with the Intel® SSD DC S3500 Series.

AppNexus found that unlocking the performance of wide SATA required specific configuration changes. Aerospike struggled with performance problems created by RAID controllers. High performance cards, like the Dell H710p* use the LSI 2208* which has the LSI Fast Path* feature that must be utilized. Without the LSI Fast Path, the LSI 2208 will bottleneck at just four SSD drives. With LSI Fast Path enabled the user can take advantage of the full speed of the wide SATA system (more SSDs), overcome the RAID bottleneck and support larger data sets. With this configuration, the Intel® SSD DC S3700 Series achieved up to 99 percent of reads within one millisecond and the Intel® SSD DC S3500 Series achieved up to 98 percent of reads in the same time frame. This configuration enables SATA-based systems with Intel® SSDs to perform at the same level as PCIe and single-level cell (SLC) SSD-based solutions at a very compelling price point.