

## PayPal Solves Fraud Challenges with Aerospike<sup>®</sup> and Intel<sup>®</sup> Optane<sup>™</sup> Persistent Memory

Deploying petabyte-scale computing with 2nd Gen Intel<sup>®</sup> Xeon<sup>®</sup> Scalable processors delivers real-time decision-making without breaking the bank



### Executive Summary

PayPal ([www.paypal.com](http://www.paypal.com)) is the world's largest online money transfer, billing, and payments system. It owns the PayPal, Venmo, iZettle, Xoom, Braintree, and Paydiant brands. By leveraging technology to make financial services and commerce more convenient, affordable, and secure, the PayPal platform empowers more than 325 million consumers and merchants in more than 200 markets to join and thrive in the global economy. But, like any banking service, PayPal faces fraud challenges. By adopting new Intel<sup>®</sup> technologies and Aerospike's real-time data platform, PayPal reduced the number of missed fraudulent transactions by 30X by improving Service Level Agreement (SLA) adherence to 99.95% up from 98.5%, while using a computing footprint 8X smaller than its previous infrastructure (1,024 servers down to 120), enabling an increase in the amount of data evaluated by 10X.<sup>1</sup>

### Challenge

PayPal's fraud rate is between 0.17 and 0.18 percent of revenue.<sup>2</sup> While this number is well below the industry average of 1.86 percent,<sup>2</sup> it still represents over 1 billion dollars a year in losses for the firm.

To identify emerging fraud patterns in real-time, PayPal needed to process and analyze more data faster. PayPal sought to quickly build a real-time decision-making platform that was highly effective while minimizing end-user friction. The platform merged big data with real-time processing, such as customer enrollment, payment, invoicing preferences, and profile data.

Scaling the amount of data processed, however, introduced challenges that included a large primary index, hundreds of petabytes of data, and a requirement to double the online transaction processing (OLTP) rate from 3.5 to 7 million transactions per second. PayPal determined that the massive amounts of data it was collecting was dragging down its existing database. The legacy database could not deliver results without a deterioration in performance. PayPal needed a cost-effective IT strategy that would allow their infrastructure to scale horizontally, while also ensuring that performance and uptime did not suffer. In 2015, they chose Aerospike and its NoSQL database system.



## Solution

PayPal's infrastructure supported over 4,000 database instances, 100 petabytes of data, and 32 percent data storage growth each year. However, as data volumes in its fraud decision-making platform grew each year, so did the size of the primary index. Once the memory capacity in a node was exhausted, it could no longer store new data, even if storage capacity was still available. Scaling out to additional nodes came at a significant cost, including upfront hardware costs, staff to manage the nodes, and increased power consumption. In order to keep database node counts down, PayPal continually looked for ways to increase the density per node.

Aerospike's database is designed with a patented Hybrid Memory Architecture™ (HMA) that delivers a fundamentally different approach from traditional databases. To deliver predictable performance up to petabyte scale, Aerospike is built explicitly to leverage the hardware advancements of flash (SSD) storage and persistent memory instead of relying on DRAM. Using Aerospike allowed PayPal to scale with higher storage density and fast SSDs, while storing only the primary index in memory for fast data access.

In 2019, PayPal began to look at next-generation memory and storage options to increase their overall storage density and reduce their DRAM costs using Aerospike. When Intel® Optane™ persistent memory (Intel® Optane™ PMem) launched, Aerospike had already optimized its software to take advantage of the new PMem modules. Intel Optane PMem delivers a unique combination of affordable large capacity and non-volatility. 2nd Gen Intel Xeon Scalable processors are designed to take advantage of Intel Optane PMem in servers.

Aerospike's HMA now offers the option to store the primary index in Intel Optane PMem instead of DRAM. Since Intel Optane PMem is available in 128, 256, or 512 GB modules at a lower cost per GB than DRAM, PayPal was able to significantly increase the memory capacity of each node. With more room for the primary index, they could now build smaller clusters with denser nodes. This greatly reduced the need to scale out, while still maintaining the high performance PayPal expects.

Persistence is an additional benefit PayPal gained by storing the primary index in Intel Optane PMem. With indexes in DRAM, a system restart required scanning data from storage to rebuild the index, but with indexes persisted across restarts, the entire system can now be restored and back online in a fraction of the time. Faster server restarts support PayPal's requirements for high uptime and better reliability.

More recently, a multiyear collaboration between Intel and Aerospike has resulted in further optimizations that include the option to store data (not just indexes) in Intel Optane PMem as well.



## Result

PayPal now has a total of 2,000 Aerospike servers, including 200 high-density servers with 2nd Gen Intel Xeon Scalable processors and Intel Optane PMem. The company expects continued growth, and by scaling its infrastructure with these new technologies, PayPal lowered TCO through DRAM reduction, while maintaining extremely fast application response and low latency. Additional benefits include:

- **4X increase** in storage per node (from 3.2 to 12 TB)<sup>1</sup>
- **2X reduction** in servers<sup>1</sup>
- **30 percent reduction** in costs per cluster<sup>1</sup>
- **12X reduction** in re-indexing time<sup>1</sup>

Based on PayPal's initial projected fraud data size of 50 TB in 2015 with its legacy in-memory system, the Aerospike HMA, Intel® processors, and Intel Optane PMem allowed PayPal to achieve the following:

- **30X reduction** in the number of fraud transactions missed by improving SLA adherence to 99.95% up from 98.5%
- **8X reduction in server footprint:** from 1,024 servers down to 120
- **Costs reduced by 3X:** USD9 million hardware cost savings projected from USD12.5 million down to USD3.5 million
- **5X throughput improvement:** from 200K to 1M transactions per second

As PayPal's annual projected data is expected to grow 32 percent, Aerospike and Intel Optane PMem enable the PayPal fraud solution to cost effectively scale. PayPal can continue to meet 99.95 percent fraud calculation SLAs. PayPal also benefits from slashed node restart times, improved query performance, and data consistency, plus maintaining consistent high availability 24/7 (99.99 percent uptime).

Fraud prevention is an important area of investment for PayPal. The company has successfully used real-time decision-making with robust fraud prevention models for more than a decade. However, fraudsters are constantly changing their patterns and uncovering new ways to take advantage of the system. As a result, PayPal must continuously find ways to improve fraud detection accuracy and decrease fraud detection time. PayPal continues to work with Aerospike and Intel to advance approaches to detect and stop fraud with higher decision confidence.

### About Aerospike®

The Aerospike Real-time Data Platform enables organizations to act instantly across billions of transactions while reducing server footprint by up to 80 percent. The Aerospike multi-cloud platform powers real-time applications with predictable sub-millisecond performance from terabytes to petabytes of data with five nines uptime with globally distributed, strongly consistent data. Applications built on the Aerospike Real-time Data Platform fight fraud, provide recommendations that dramatically increase shopping cart size, enable global digital payments, and deliver hyper-personalized user experiences to tens of millions of customers. Customers such as Airtel, Experian, Nielsen, PayPal, Snap, Verizon Media, and Wayfair rely on Aerospike as their data foundation for the future. Headquartered in Mountain View, California, the company also has offices in London, Bangalore, and Tel Aviv.



<sup>1</sup>Testing was conducted by Aerospike and PayPal at PayPal in April 2020. 4x increase in storage: The baseline configuration has 6.4 TB/node, while the new configuration with Intel Optane persistent memory has 25.6 TB/node (6.4 X 4 = 25.6). 2x reduction in servers: The baseline cluster had 20 nodes holding 2 copies of each record, while the new cluster with Intel Optane persistent memory has only 10 nodes that hold 3 copies of each record, enabling a higher level of fault tolerance. 30% reduction in cost/cluster: Baseline cluster of 20 nodes costs \$1.3 million where each node costs about \$65K to the new cluster of 10 nodes at a cost of \$0.9 million, each node costs about \$90K (\$0.9 million / \$1.3 million = 0.70, so 30% reduction in cluster cost). 12x reduction in re-indexing time: From 59 minutes to re-index on the original configuration to only 4 minutes on the new config with Intel Optane persistent memory.

<sup>2</sup>LexisNexis® Risk Solutions 2019 True Cost of Fraud Study. <https://risk.lexisnexis.com/insights-resources/research/2019-true-cost-of-fraud-study-e-commerce-retail-edition>.

#### Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](http://www.intel.com/PerformanceIndex).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.