

SOLUTION BRIEF

Big Data / Analytics

Financial Services, AdTech, MarTech,
Telecommunications, and More



Accelerated Real-Time Computing at Petabyte Scale

Aerospike Enterprise Edition 4.8 is engineered to make advanced use of Intel® Optane™ persistent memory for real-time computing. A multiyear collaboration between the two companies enables optimizations for higher performance than SSDs, lower cost per GB than DRAM,¹ and data structure storage that persists over system restarts.

Ten years ago, Aerospike database helped change the game for real-time computing by providing the performance and scale needed to support high volume, mission-critical transactions and analytics on a single data platform. Leading companies in advertising, retail, financial services, telecommunications, and many other industries² have taken advantage of these capabilities to integrate smarter, faster decision making into some of their most critical, high volume, and time-sensitive business processes.

Now Aerospike, working with Intel, has changed the game once again. With deep optimizations for Intel Optane persistent memory, Aerospike Enterprise Edition 4.8 offers unprecedented scale for real-time computing, at affordable cost, and with even higher availability than dynamic random-access memory (DRAM) only solutions (Figure 1).

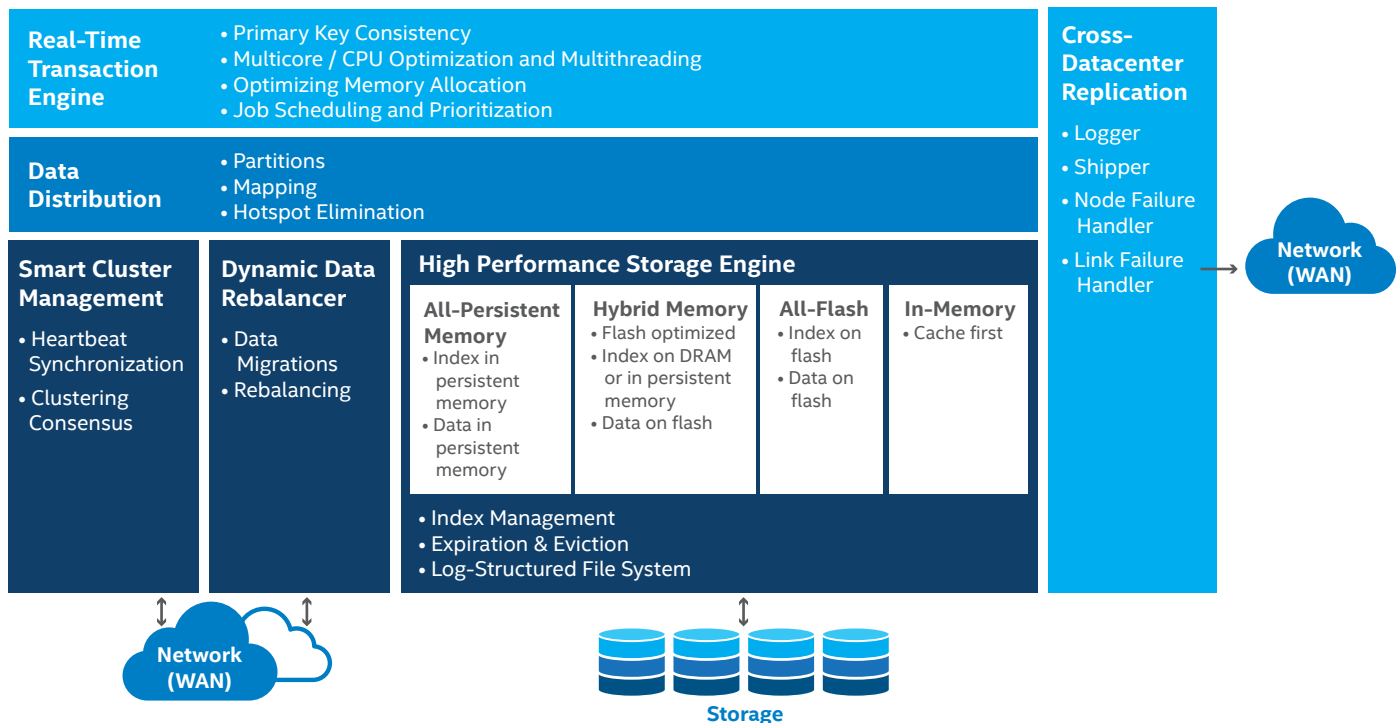


Figure 1. Aerospike 4.8 with Intel Optane persistent memory supports high volume, real time, mission-critical transactions and analytics operating on petabyte-scale data sets with mission-critical reliability.

Data-hungry applications, such as fraud detection, digital payment systems, real-time bidding, and recommendation engines can now tap into petabytes of data to support real-time processes. Businesses no longer must limit the amount of data they analyze in order to meet extreme performance requirements. They can extract deeper and more reliable insights. They can also integrate new machine learning and deep learning algorithms more easily into time-sensitive business processes.

In addition to enabling more memory and larger data volumes per server, Aerospike 4.8 running on a 2nd Gen Intel Xeon Scalable processor-based platform equipped with Intel Optane persistent memory also simplifies and improves high availability. Data and database indexes can be retained in persistent memory when the system is powered down, so Aerospike can typically be restarted in a matter of seconds to enable non-disruptive maintenance. Downtime is reduced, software updates and security patches can be performed more frequently, and redundancy and replication requirements can be met more easily, all at potentially lower cost. In combination with Aerospike's automated clustering, cross data center replication, and strong consistency, the value proposition for real time, mission-critical computing has never been stronger.

Breakthrough Memory Architecture for Data-Hungry Applications

As a distributed enterprise database, Aerospike enables near-instant access³ to large data sets for real-time computing. Clustered server architectures use the cumulative memory footprints of many nodes to achieve high performance at scale. Although this approach scales efficiently, upper limits have been constrained by the high cost and limited capacity of DRAM.

Resolving the Memory Bottleneck for Faster, Deeper Insights

Intel Optane persistent memory pushes through those cost and capacity constraints with modules available in 128 gigabytes (GB), 256 GB, and 512 GB capacities in a form factor that is pin compatible with conventional DRAM. Because this technology offers high performance, byte addressable memory at a significantly lower cost per GB than DRAM,¹ it enables much more memory per server than is possible with today's widely used 32 GB DRAM DIMMs. (While 64 GB and 128 GB DRAM modules are also available, they typically cost more per GB, making cost even more of a stumbling block.)

With proper enablement, software can also take advantage of the non-volatility of Intel Optane persistent memory. Unlike with DRAM, stored data is retained when the power goes down, providing more rapid recovery from system restarts for a more stable and resilient enterprise data foundation.

Verizon Builds a More Immersive Future with Data

Building out its Global Data Platform to power a new generation of capabilities for personalization and serving customer data, Verizon Media is streamlining and standardizing a broad portfolio of database infrastructures onto Aerospike, enabling advanced user profile and data aggregation solutions. The new infrastructure will be powered by hundreds of 2nd Gen Intel Xeon Gold processors, as well as more than 420 terabytes (TB) of Intel Optane persistent memory.

More than two years of co-engineering and enablement by Aerospike and Intel have culminated in the availability of this cutting edge database solution, and the marketplace is taking notice. Similar solution stacks based on technology building blocks from Aerospike and Intel are entering production across industry segments by companies as diverse as PayPal, Wayfair, The Trade Desk, and AirTel.



This new technology redefines the memory tier, offering advantages as shown in Figure 2.



Figure 2. Technology and business advantages of Intel® Optane™ persistent memory.

Traditionally, the amount of system memory available constrained how many primary index keys could be hosted per server node, which in turn limited how much user data could be hosted per node. Intel Optane persistent memory lifts these constraints. With the introduction of Intel Optane persistent memory, increasing memory capacity became practical as well as cost-effective. Each node of an Aerospike cluster can potentially support up to 4X as much data as a DRAM only solution, as illustrated in Figure 3.⁴ An Aerospike cluster can be scaled to 128 nodes to handle large datasets and heavy transaction volumes, without increasing latency.




Figure 3. Intel Optane persistent memory enables memory scalability by up to 4X versus DRAM-only solutions.⁴

Enabling Flexibility with Persistent Memory Operating Modes

Intel Optane persistent memory offers two types of operation: Memory Mode and App Direct Mode. This flexibility enables software providers with a simple path to adoption as well as the potential for deep optimizations to drive very high performance at petabyte scale. Aerospike Enterprise Edition 4.8 exemplifies deep commitment by a software provider to engage in ongoing tuning and optimizations for App Direct Mode, to drive the fullest value possible from Intel Optane persistent memory.

- **Memory Mode** extends the memory you have and uses DDR4 memory to act as a high-speed data cache controlled by the processor's memory controller. It does not take advantage of memory persistence. No specific optimizations or other changes are required to software applications, and the data will not be saved in the event of a power loss.⁵

- **App Direct Mode** delivers a tiered approach—one volatile and one persistent that gives applications the option to optimize where data is placed. It requires modification of application software, which enables substantially increased performance. The application and operating system are explicitly aware there are two types of direct load/store memory and have the ability to write directly to persistent memory. Accessing data directly from the memory bus delivers response times measured in nanoseconds, compared to milliseconds for the I/O bus.⁶

Ongoing, Advanced Enablement in Aerospike 4.8

Aerospike 4.8 provides a new, higher level of optimization for Intel Optane persistent memory than predecessor versions of the database, further improving the processor's access to data to increase performance and scale. Aerospike has remained committed to taking advantage of the increasing availability of high-speed storage, resetting existing price/performance expectations, and multiple deployment options, as illustrated in Figure 4.

Typical implementations of Aerospike until recently have been deployed on DRAM only servers, holding the database index (the keys) in high-speed DRAM, and the associated data (the values) in high-capacity SSDs. This approach was necessarily constrained by the compromise between the superior speed of DRAM and the lower cost per gigabyte of SSDs. Because the index was held in volatile DRAM, it would need to be rebuilt each time the server was restarted, a process that potentially took hours for petabyte-scale data.

Aerospike 4.5 and later versions are optimized for Intel Optane persistent memory App Direct Mode, with higher memory capacities than DRAM, enabling larger indexes in addition to the broader set of performance advantages that App Direct Mode provides. Intel Optane persistent memory provides lower cost per gigabyte than DRAM and lower retrieval latency than SSDs. At the same time, the index is persistent over power cycles, meaning that it does not need to be rebuilt at system reboot.

Beginning with Aerospike 4.8, data as well as database indexes can be stored in Intel Optane persistent memory. The memory tier therefore functions as the main data capacity tier, with the main data store housed in Intel Optane persistent memory. DRAM still acts as a high-speed cache for hot data, with lower latency access to warm and cool data than with SSDs. This architecture provides very high data throughput with very low latency for the most performance-critical real-time use cases. Applications make direct loads and stores to the media, at cache-line granularity.

Note: Customers must ensure that each server node has sufficient processing and network capacity to handle the increased transaction loads associated with higher data volumes. Performance testing with actual workloads is recommended to determine the optimal processor, memory, and storage configurations.

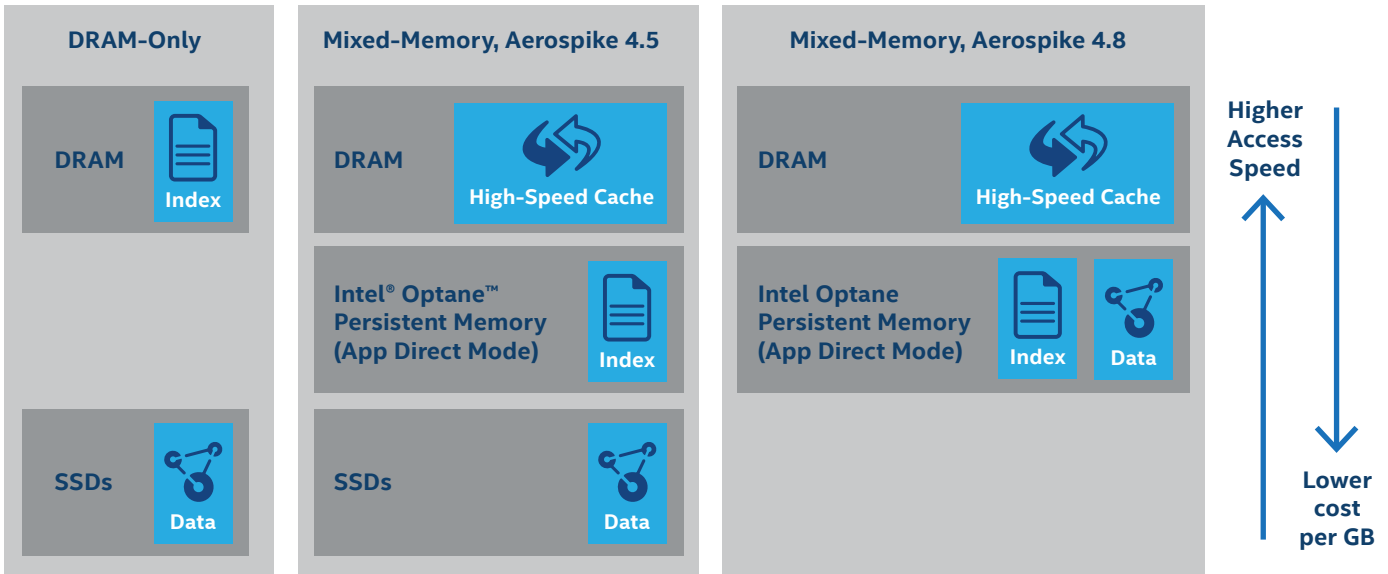


Figure 4. Evolution of Aerospike and Intel collaboration.

On Premises or in the Cloud

Prior to the launch of Intel Optane persistent memory, Intel worked extensively with commercial and open source leaders to pave the way for broad adoption. Aerospike Enterprise Edition was the first commercially available open database to provide optimized support for App Direct Mode and the persistence it enables. A number of leading hardware, OS, hypervisor, and application vendors have also been preparing for the release of Intel Optane persistent memory, and many have deployment-ready solutions available today (for details, read [Data-Centric Innovation Spotlight Series, Big Memory Breakthrough for Your Biggest Data Challenges](#)).

Note: The optimizations made by Aerospike for Intel Optane persistent memory are distinct from Storage over App Direct Mode, a separate approach to enablement taken by some software makers. This approach does not require software modification, and accesses the persistent memory in 4 KB or larger blocks using a standard file API, treating it essentially as a storage device. While performance using Storage over App Direct Mode is better than with SSDs, it is substantially lower than with App Direct optimization.

Refer to <http://software.intel.com/en-us/articles/optimization-notice> for more information regarding performance and optimization choices in Intel software products.

Optimizing Software for Intel Optane Persistent Memory

For software providers to unleash the benefits of App Direct Mode, Intel has developed the **Persistent Memory Development Kit (PMDK)** to simplify and accelerate the required software modifications. The PMDK is open source and based on the Storage Networking Industry Association (SNIA) programming model. It is freely available and offers multiple levels of granularity. Because Aerospike is focused on delivering high performance at large scale, their engineering team chose to make strategic changes to their code base at the most granular level. Other organizations may be looking for quicker time to value and may choose a less granular approach.

Intel offers additional information, training, and tools that can help developers understand and evaluate Intel Optane persistent memory for particular use cases. These resources include tools and methods for analyzing existing applications to determine if big, affordable, persistent memory can help to achieve lower latency, higher resilience, and greater performance. For more information, visit the **Intel® Developer Zone**.

Unleash Your Data, Free Your Business

Aerospike 4.8 and Intel Optane persistent memory are changing the game for smart, real-time business solutions. It is now possible—and practical—to take advantage of petabyte-scale data sets for transactions and analytics that operate at both high volume (millions of transactions per second) and low latency (milliseconds). With this speed and scale, organizations can analyze newly generated data in combination with historical data using deep learning and other analytical approaches. Most importantly, they can do it fast enough to make better decisions during real-time mission-critical interactions.

Aerospike is focused on being the data platform of choice for the next generation of real-time business applications. They are accomplishing this by supporting geographically distributed data sets across both private data centers and public clouds as a system of record, while also enabling high-speed edge processing, so businesses can embed deep intelligence into their real-time business models. In combination with 2nd Gen Intel Xeon Scalable processor-based platforms and Intel Optane persistent memory, Aerospike 4.8 is delivering on those requirements today with greater scale, higher availability, and better cost models than ever before.


Learn more about:

Aerospike database: www.aerospike.com

Intel® Optane™ Persistent Memory: www.intel.com/OptanePersistentMemory

Intel® Developer Zone: <http://software.intel.com/pmem>

Solution provided by:



¹ Intel Optane persistent memory pricing & DRAM pricing referenced in TCO calculations is provided for guidance and planning purposes only and does not constitute a final offer. Pricing guidance is subject to change and may revise up or down based on market dynamics. Please contact your OEM/distributor for actual pricing. Your cost and results may vary.

² For dozens of Aerospike customers case studies, visit <https://www.aerospike.com/customers/>.

³ Tests performed by Intel and Aerospike as of 27 February 2019 demonstrated that a server configured with DRAM + Intel Optane persistent memory provided 95.74 percent of the performance of a system configured with an equivalent amount of DRAM only memory.

Baseline configuration: Intel® Xeon® Platinum 8280 processor (2.7 GHz, 28 cores), 1.5 TB total memory (24 x 64 GB @ 2666 MT/s DDR4 LRDIMM), 1 x Intel SSD DC S3700 (800 GB) + 7 x Intel® SSD P4510 (2 TB) 2.5" PCIe, CentOS Linux 7.4 kernel 4.19.8.

New configuration: Intel Xeon Platinum 8280 processor (2.7 GHz, 28 cores), 192 GB total memory, 12 slots x 16 GB @ 2666 MT/s DDR4 RDIMM plus 12 slots x 128 GB Intel® Optane™ persistent memory, 1 x Intel SSD DC S3700 (800 GB) plus 7 x Intel SSD P4510 (2 TB) 2.5" PCIe, CentOS Linux 7.4 kernel 4.19.8.

⁴ The up to 4X increase in memory per server is based on a typical, two-socket server configured with 24 x 64 GB DRAM DIMMs (= 1.5 TB of memory) versus the same server configured with 12 x 512 GB Intel® Optane™ persistent memory modules (= 6.0 TB of memory).

⁵ Aerospike has not tested Intel Optane persistent memory in Memory Mode.

⁶ Testing by Intel as of February 22, 2019. Latency versus load, 70Read/30Write Random, 4 KB for SSD and 256 B for memory.

SSD system: 2x Intel® Xeon® Gold 6154 processor @ 3.0 GHz (18 cores), 256 GB DDR4-2666, Intel® Optane™ SSD DC P4800X (375 GB), Intel® SSD DC P4610 (3.2 TB), CentOS 7.5, kernel 4.17.6-1.el7.x86_64. Intel Microcode: 0x2000043; System BIOS: 00.01.0013; ME Firmware: 04.00.04.294; BMC Firmware: 1.43.91f76955; FRUSDR: 1.43.

Persistent memory system: 1x Intel® Xeon® Platinum 8276 processor @ 2.2 GHz (28 cores), 32 GB DDR4-2666, Intel® Optane™ persistent memory (256 GB). AEP FW: 5336, BIOS: 573.D10, BKC version: WW08 BKC, Linux OS: 4.20.4-200.fc29, Spectre/Meltdown Patched (1,2,3,3a), Performance Turning: QoS Disabled, IODC=5(AD).

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Performance results may not reflect all publicly available security updates. See configuration disclosure for details.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance/datacenter>.

Intel compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

© 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.

0420/KC/MESH/PDF 338809-003US