SOLUTION BRIEF

Big Data / Analytics
Financial Services, AdTech, MarTech, Telecommunications, and More

Affordable Real-Time Computing at Petabyte Scale

Aerospike® Enterprise Edition 4.5, together with Intel® Optane™ DC persistent memory, provides a breakthrough in scalability, high availability, and affordability for real-time computing.

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 industries[1] 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 DC persistent memory, Aerospike Enterprise Edition 4.5 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).

In practical terms, this means that 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 have to 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 4.5 with Intel Optane DC persistent memory also simplifies and improves high availability. 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.

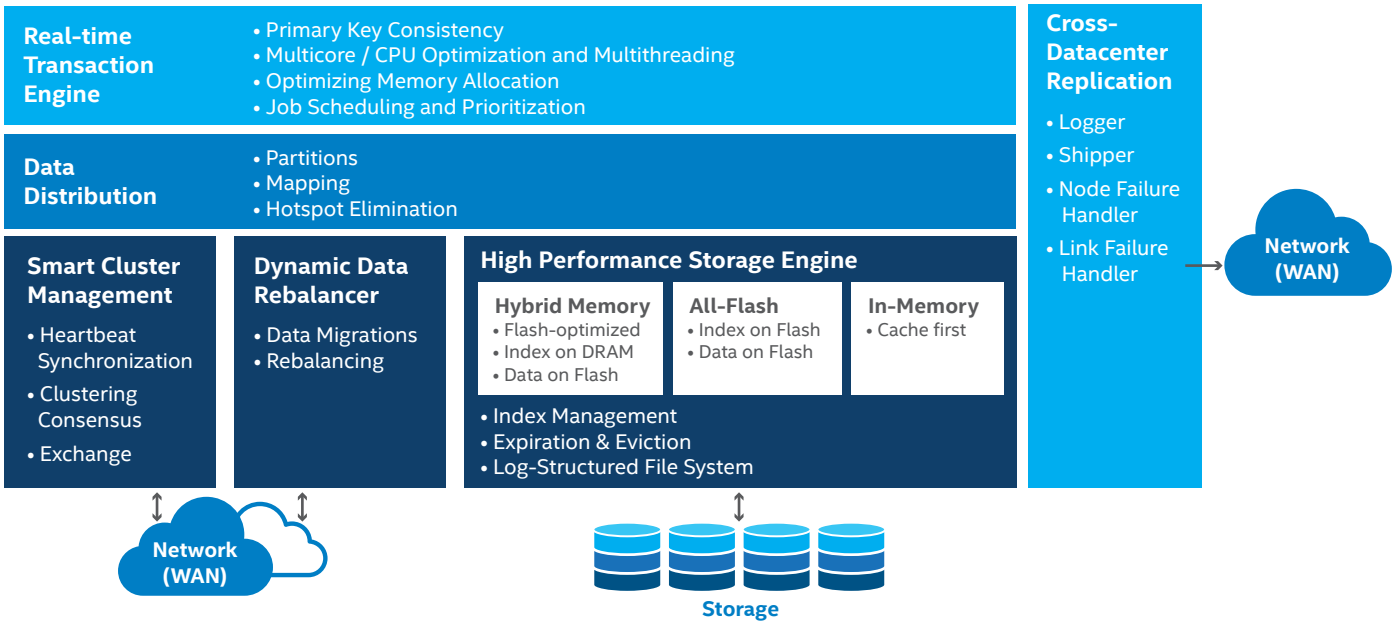

Figure 1. Aerospike 4.5 with Intel® Optane™ DC persistent memory supports high-volume, real-time, mission-critical transactions and analytics operating on petabyte-scale data sets with mission-critical reliability.

A Breakthrough for Data Hungry Applications

Distributed, flash-optimized databases such as Aerospike enable near-instant access to large data sets for real-time computing. They do this with clustered server architectures that use the cumulative memory footprints of many nodes to achieve high performance at scale. Although this approach scales with near-linear efficiency, upper limits have been constrained by the high cost and limited capacity of traditional DRAM memory.

Resolving the Memory Bottleneck for Faster, Deeper Insights

Intel Optane DC persistent memory pushes beyond the traditional constraints of DRAM. Each module provides 128 gigabyte (GB), 256 GB, or 512 GB of affordable, high-performance, byte-addressable memory in a standard dual in-line memory module (DIMM) form factor. This enables much more memory per server than is possible with today's widely used 32 GB DRAM DIMMs (larger 64 GB and 128 GB DIMMs are available, but typically cost more per GB, which makes them cost prohibitive for many use cases). Intel Optane DC persistent memory is also non-volatile. Unlike with DRAM, stored data is retained if the power goes down, providing a more stable and resilient foundation for enterprise data.

In other words, this new technology combines the best of memory and storage in a way that changes both what is practical and what is possible in today's data center and cloud environments. Not only is maximum memory capacity greatly expanded, but persistent data storage can potentially be accessed orders of magnitude faster than from even the fastest solid-state drives (SSDs). This unprecedented combination of performance, capacity, persistence, and affordability breaks down long-standing barriers to high-speed, high-volume data access. It also helps to reduce the need for high-volume data movement, which can significantly ease the load on network infrastructure.

Two Modes of Operation

Intel Optane DC persistent memory can be used in two distinct modes.

- In **Memory Mode**, Intel Optane DC persistent memory operates as volatile memory for the server platform, and a small amount of DRAM provides a cache for the most frequently used data. In essence, the amount of system memory can be greatly expanded, with little to no impact on memory performance.
- **App Direct Mode**, which is enabled in Aerospike Enterprise Edition 4.5, takes advantage of the inherent persistence of Intel Optane DC persistent memory. The application sees two distinct memory resources: one volatile and one persistent. It sees traditional DRAM as fast, volatile memory and Intel Optane DC persistent memory as high-capacity persistent memory. Although software modification is required at the application level, App Direct Mode allows the application to control data placement. In many cases, this can help to improve performance, since developers can take advantage of the different strengths of DRAM and Intel Optane DC persistent memory.

Deep Integration in Aerospike 4.5

Aerospike is a key value database. Its fast, predictable performance is largely attributable to the fact that it typically holds the database index (the keys) in high-speed DRAM memory, and the associated data (the values) in high-performance SSDs.

With this approach, data is always accessed in the same way, which ensures that latency is low and predictable. An Aerospike cluster can currently be scaled to 128 nodes to handle large data sets and heavy transaction volumes, all without increasing latency. Until now, the amount of DRAM in the system constrained how many keys could be hosted per server node, which in turn limited how much user data could be hosted per node. Intel Optane DC persistent memory lifts these constraints to provide several benefits.

A 4X Increase in Scalability[2]

When using only 64 GB DRAM DIMMs for system memory, a single two-socket server node is limited to about 1.5 terabytes (TB). With Aerospike, this is enough memory to hold the index for roughly 22.5 billion records. Using Intel Optane DC persistent memory in combination with a small amount of DRAM, a two-socket server node can readily be configured with up to 6.0 TB of memory, enough to support 90 billion records.[2] (The DRAM in the mixed-memory system can be used for either secondary indexes or hot data.)

In other words, with Intel Optane DC persistent memory, each node of an Aerospike cluster can potentially support up to 4X as much data as a DRAM-only solution (Figure 2).

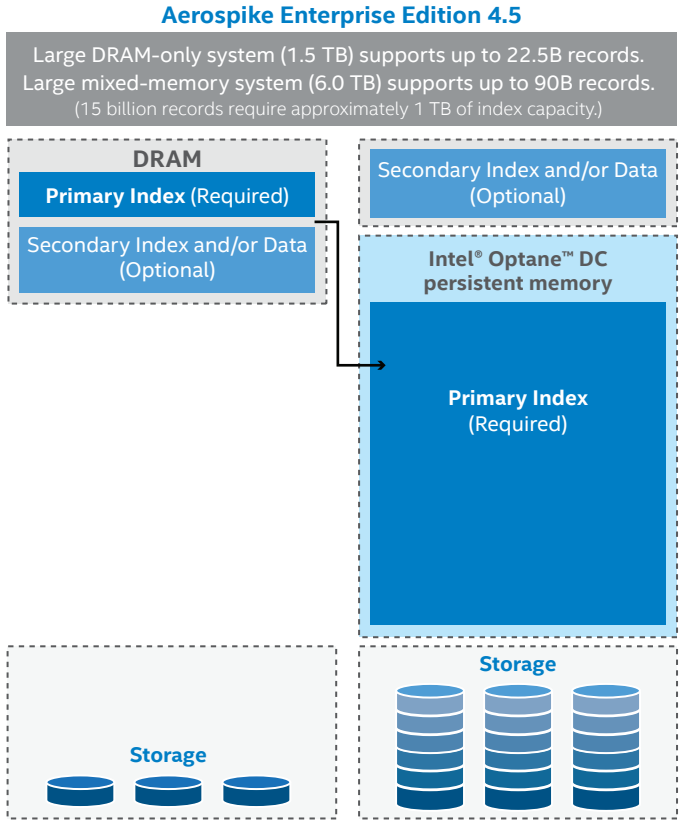

Figure 2. Aerospike 4.5 with Intel® Optane™ DC persistent memory increases scalability by up to 4X versus DRAM-only solutions, while improving reliability and providing the same extreme levels of performance.

Blazingly Fast Performance

Aerospike 4.5 is tuned to make efficient use of all the memory resources in each server node, including both DRAM and Intel Optane DC persistent memory. Data placement is optimized to take advantage of the slightly faster performance of DRAM and the higher capacity and persistence of Intel Optane DC persistent memory.

Based on Aerospike benchmark data, performance for a system with Intel Optane DC persistent memory and a relatively small amount of DRAM is only slightly less than the performance of a system configured with an equivalent amount of DRAM-only memory.[3] Combining the two memory types can, therefore, provide the benefits of persistence, at potentially lower cost and with sufficient performance for a wide range of real-time use cases.

Alternatively, Intel Optane DC persistent memory can be used to reduce the number of required servers in an Aerospike 4.5 cluster, since increasing the amount of total memory per server will typically allow larger storage configurations per server without sacrificing performance. This can significantly reduce both capital and operating costs. However, customers will need to be sure they have sufficient processing capacity in each server node to handle the increased transaction loads associated with the higher data volume. Performance testing with actual workloads is recommended to determine the optimal processor, memory, and storage configurations.

Higher Uptime

With a traditional DRAM-only Aerospike solution, the database index is lost whenever the system is powered down for maintenance or upgrades. User data must then be scanned during startup to rebuild the indexes, a process that can take a considerable amount of time; typically several hours.

With Aerospike 4.5 and Intel Optane DC persistent memory, the entire database index can be retained in persistent memory when power is cycled off for planned downtime. With this approach, database startup times can be reduced by as much as 135 times[4] for planned reboots. A startup can be completed in a matter of seconds rather than minutes or hours, so maintenance and security updates can be conducted with very little disruption to production systems. This can be a major advantage for databases that are used to support mission-critical, time-sensitive applications.

On Premises or in the Cloud

Prior to the launch of Intel Optane DC persistent memory, Intel worked extensively with commercial and open source leaders to pave the way for broad adoption. Aerospike was one of the first vendors to show interest, and Aerospike Enterprise Edition 4.5 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, operating system (OS), hypervisor, and application vendors have also been preparing for the release of Intel Optane DC 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*).

Perhaps most importantly, several cloud service providers are working to provide large-memory server instances that include Intel Optane DC persistent memory. With these solutions, IT organizations will be able to have Aerospike 4.5 instances up and running quickly, so they can upload their data and begin testing performance using their own workloads. A proof-of-concept can be run quickly and at low cost.

Optimizing Software for Intel® Optane DC™ Persistent Memory

Intel Optane DC persistent memory offers value for a wide range of data-intensive workloads that are currently constrained by the limitations of traditional DRAM memory. Users and software vendors can take advantage of Memory Mode to increase memory capacity per server and to potentially reduce the cost of larger memory configurations. No code modifications to the application are required, although a persistence-aware OS is required, and some code optimization may be useful to take full advantage of the expanded memory footprint.

For those interested in 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 DC 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 4.5 and Intel Optane DC 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 Intel Optane DC persistent memory, Aerospike 4.5 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™ DC Persistent Memory: **www.intel.com/OptaneDCpersistentmemory**

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

Solution provided by:

AEROSPIKE | intel

[1] For dozens of Aerospike customers case studies, visit https://www.aerospike.com/customers/.

[2] The 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™ DC persistent memory modules (= 6.0 TB of memory).

[3] Tests performed by Intel and Aerospike as of 27 February 2019 demonstrated that a server configured with DRAM + Intel® Optane™ DC 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 processor 8280 (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 processor 8280 (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™ DC 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] Tests performed by Intel and Aerospike as of 27 February 2019 demonstrated that an 8.7 TB database could be restarted in as little as 35 seconds on a server configured with DRAM + Intel® Optane™ DC persistent memory versus up to 4,745 seconds (1 hr 19 m 5 s) on a server configured with an equivalent amount of DRAM-only memory.
Baseline configuration: Intel® Xeon® Platinum processor 8280 (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 (2TB) 2.5" PCIe, CentOS Linux* 7.4 kernel 4.19.8.
New Configuration: Intel® Xeon® Platinum processor 8280 (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™ DC 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.