

Accelerate business applications
with optimized database
performance



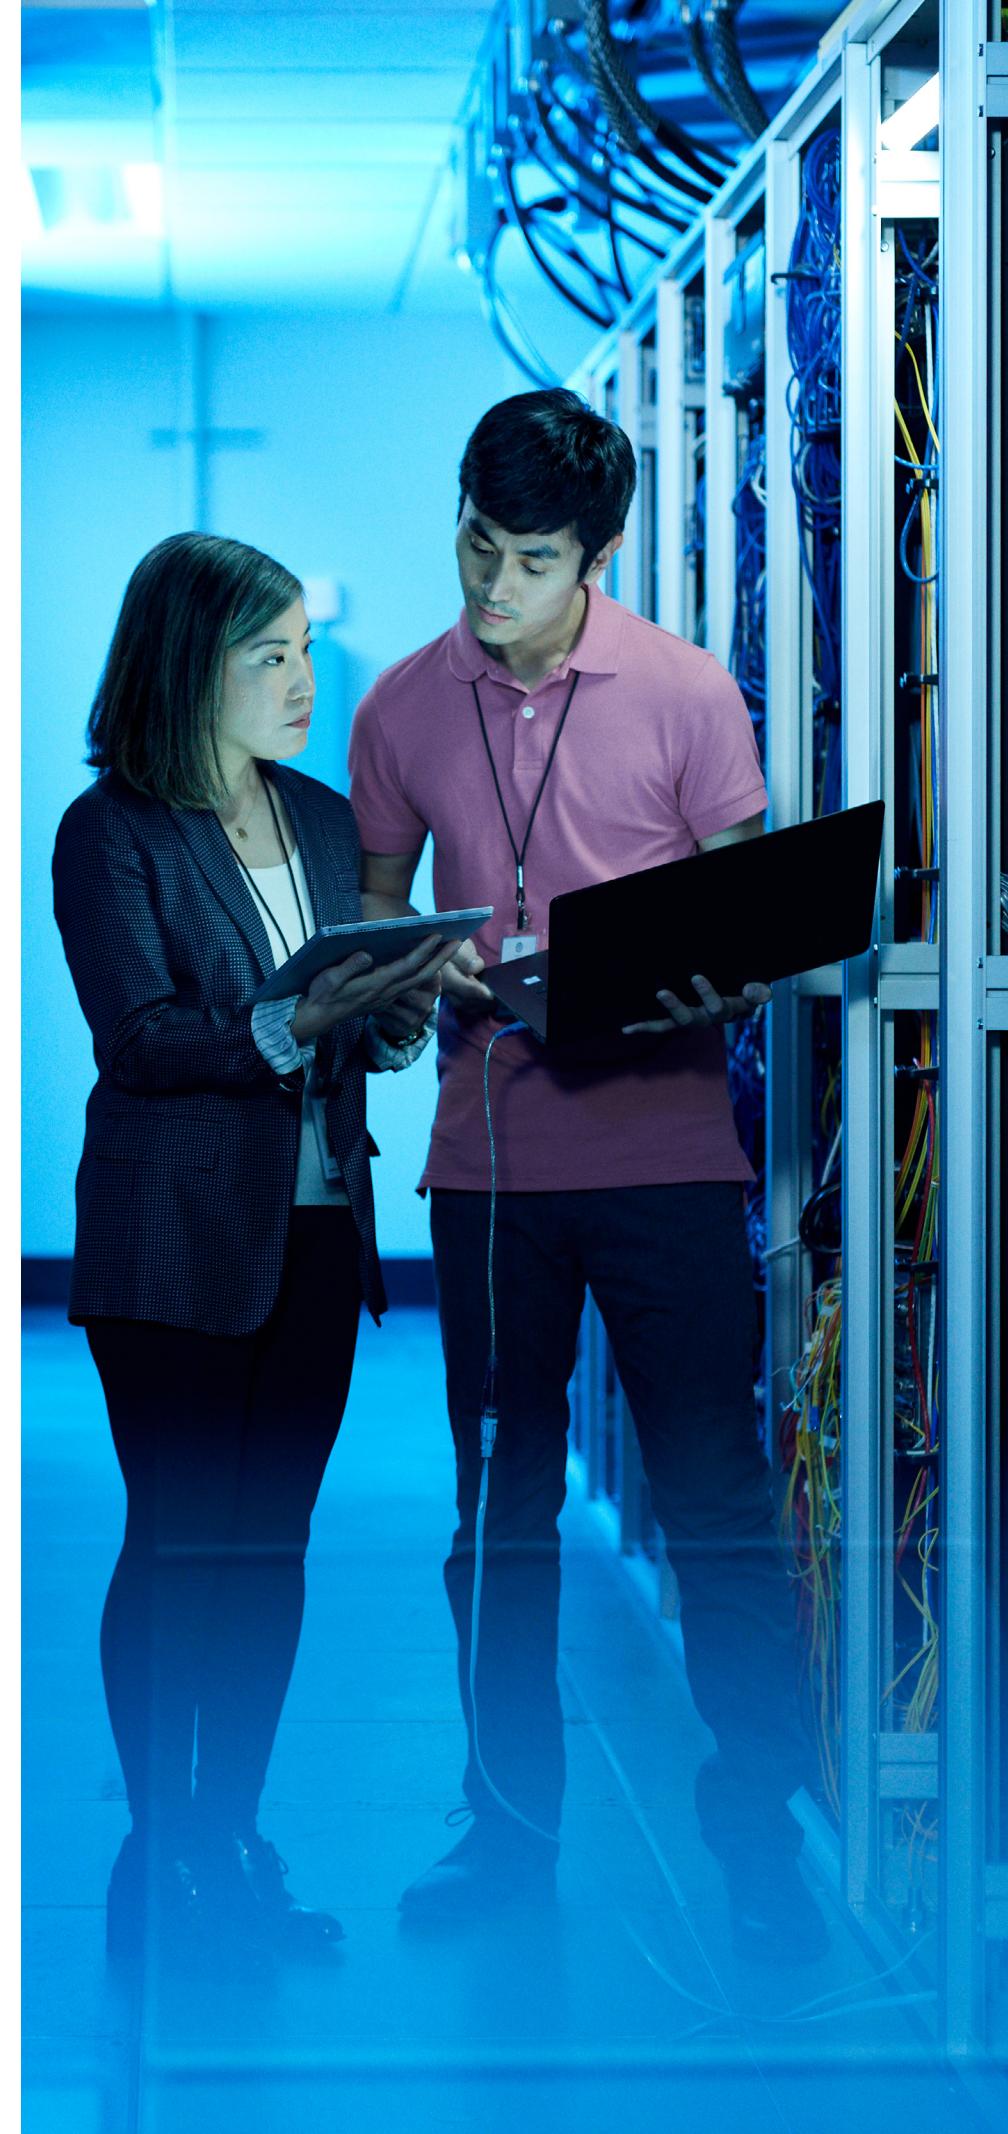
Contents

Your applications need lots of data, fast	2
Memory and storage: completing the pyramid	3
Traditional tiers of memory and storage	3
A new memory tier	4
Additional benefits of the 2nd generation Intel® Xeon® Scalable processor	4
Improving caching with Memory Mode	5
More responsive storage with App Direct Mode	6
Crawl, walk, run	8
The role of SSDs	8
Acceleration technologies	9
Deploying FPGAs	9
Apache Cassandra	10
Apache Spark	10
MySQL	10
Intel® Cache Acceleration Technology	10
Boosting connectivity	11
Conclusion	11
Learn more	12

Your applications need lots of data, fast

Business applications from human resources (HR) portals and customer relationship management (CRM) systems, to programs using advanced analytics and artificial intelligence (AI) to drive real-time insights all need fast access to data. Companies use these tools to increase their productivity, so delays will affect their businesses.

As data volumes grow, enterprises need to ensure their databases not only have the capacity for high volumes of data, but can also access it at speed.



Memory and storage: completing the pyramid

Traditional tiers of memory and storage

Building an infrastructure with the right storage solution is the foundation for an efficient and optimized database. Traditionally models of storage and memory can be thought of as a pyramid, illustrated in Figure 1. The width of the pyramid represents storage capacity, with hard drives or tape providing the largest capacity; and the height of the pyramid represents performance, with DRAM providing the fastest performance with a latency of less than 30 nanoseconds.

Striking the right balance

The model shows that infrastructure architects face a delicate balancing act. Unable to get low latency and high capacity in the same unit, they have to find the right combination of storage and memory to meet their needs. DRAM is fast but it's relatively expensive, and capacity is low. As you move down the pyramid, you find Intel® Optane™ DC SSDs, which can store up to 750GB per drive, and Intel® QLC 3D NAND SSDs, which provide terabytes of storage, with a latency of hundreds of microseconds. This makes them slower than DRAM but faster than hard disks.

Now, there is a new option which can be used for workloads that demand both speed and capacity: Intel® Optane™ DC persistent memory.

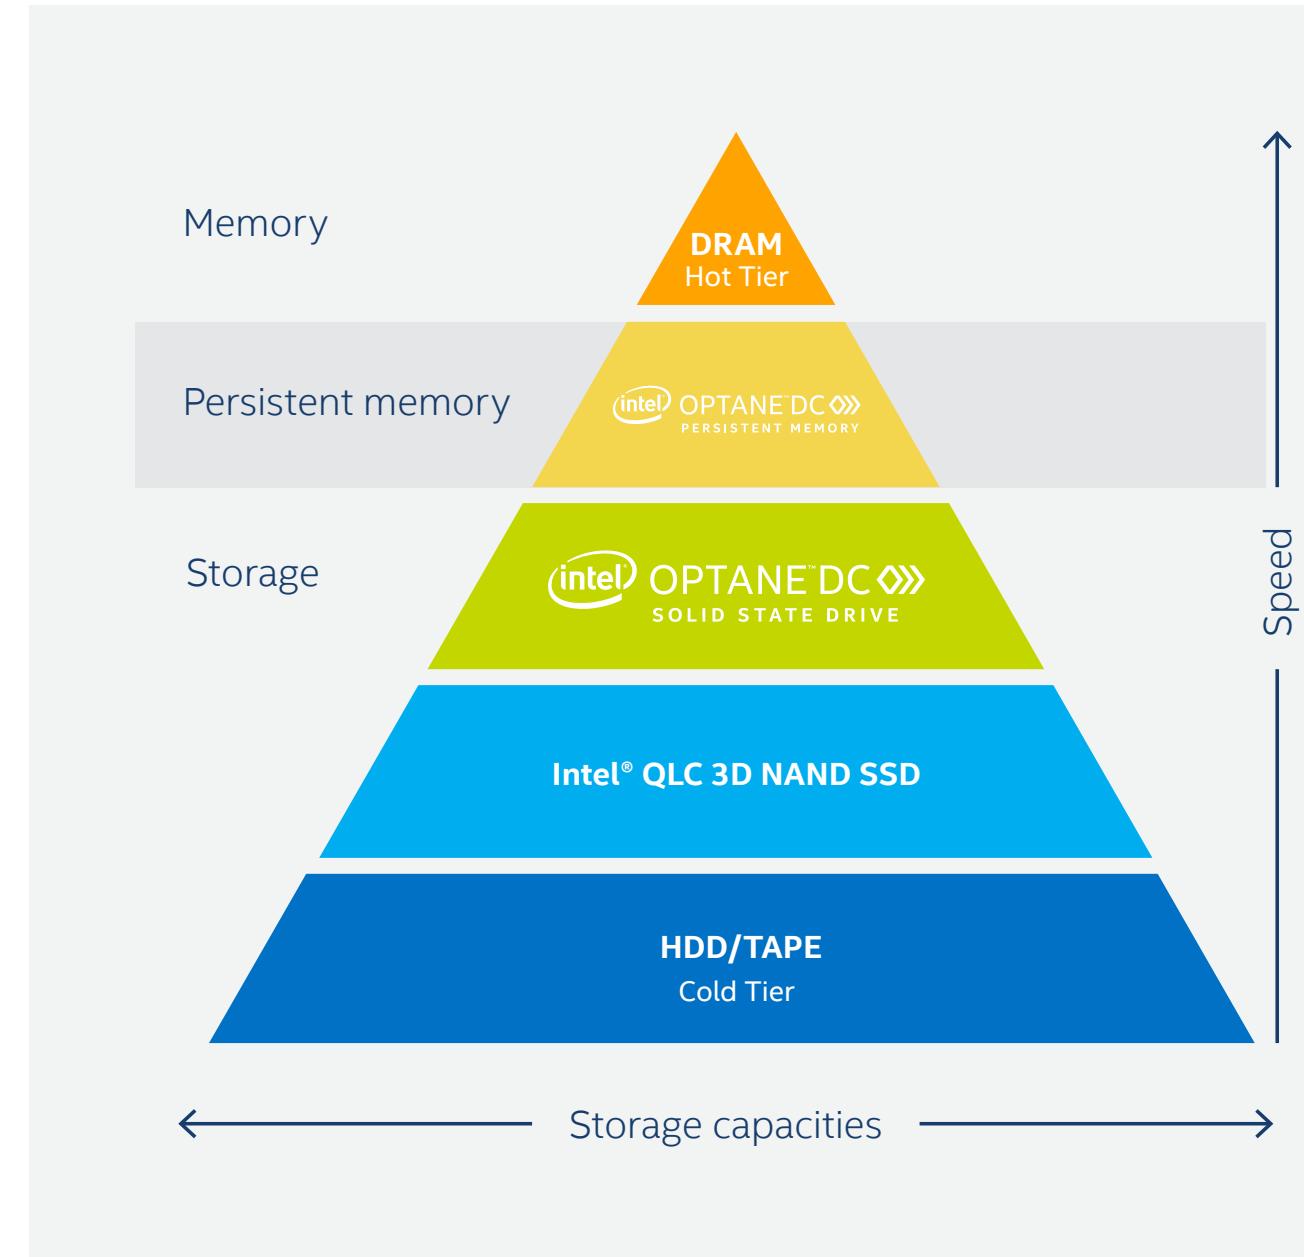


Figure 1: Storage pyramid, showing where the new Intel® Optane™ DC persistent memory sits in the storage hierarchy.

A new memory tier

The launch of the 2nd generation Intel® Xeon® Scalable processor introduced support for Intel Optane DC persistent memory. This has unlocked a new tier in the pyramid, which can help organizations access the levels of speed and capacity they need to optimize their database performance. It combines the speed of memory with affordable large capacities.

Although Intel Optane DC persistent memory is slightly slower than DRAM, it offers greater capacity than DRAM at a lower cost per bit; and greater throughput than SSDs, depending on the use case.

Intel Optane DC persistent memory operates in two modes, which can be used in different ways to improve database performance. These are called Memory Mode and App Direct Mode. Memory Mode can be used to lower the cost of the caching tier, and App Direct Mode provides high speed bulk data storage.

Additional benefits of the 2nd generation Intel® Xeon® Scalable processor

The 2nd generation Intel Xeon Scalable processor family offers up to 56 cores for high performance in compute-intensive workloads. It also supports integrated Intel® QuickAssist Technology (Intel® QAT), which can accelerate storage with real-time data compression, static data compression, and security features. It can be used to help encrypt network traffic between servers, for example, to improve database security.

The 2nd generation Intel Xeon Scalable processor is also equipped with Intel® Advanced Vector Extensions 512 (Intel® AVX-512), which can process twice as many floating-point operations per second compared to the previous-generation Intel® Advanced Vector Extensions 2 (Intel® AVX2)¹. Databases that take advantage of Intel AVX-512 can use it to help accelerate their encryption and compression.



Improving caching with Memory Mode

Large scale-out databases that run across a lot of nodes, such as Cassandra or MongoDB, can get so large that access times end up suffering, as it takes longer to find and access the right data. Using a cache in DRAM can help to accelerate access to the most frequently used data. A few years ago, you might have expected to see systems using 384GB of DRAM per node. However, it's now possible to see a cache node measured at 1.5TB. As DRAM is expensive, the need for growing cache sizes can eat into the profits of B2B SaaS providers. This is where the affordability of Intel Optane DC persistent memory can offer a real advantage.

In Memory Mode, Intel Optane DC persistent memory can be used to provide more capacity to the database cache. In this mode, it works in the same way as DRAM, with a slightly lower performance but at a reduced cost per bit.

When it comes to deployment, a key advantage of using Intel Optane DC persistent memory in Memory Mode is that it doesn't require changes to applications or operating systems.

As long as you're running on 2nd generation Intel Xeon Scalable processors, you can plug Intel Optane DC persistent memory into the DIMM slots. All the DRAM in the system is used as cache for the persistent memory. Depending on the application you are using, the ratio of DRAM to Intel Optane DC persistent memory will vary. However, a ratio of 1:8 is likely to work well for many applications.

Persistent memory should be used evenly across the platform to make better use of memory bandwidth, making persistent memory ideal for capacities between 512GB and 6TB per two-socket platform.

Redis and Memcached are often used as cache in front of databases such as MongoDB and Cassandra. Redis and Memcached both work well with Intel Optane DC persistent memory in Memory Mode (see Figures 2 and 3). These results were achieved without code changes.

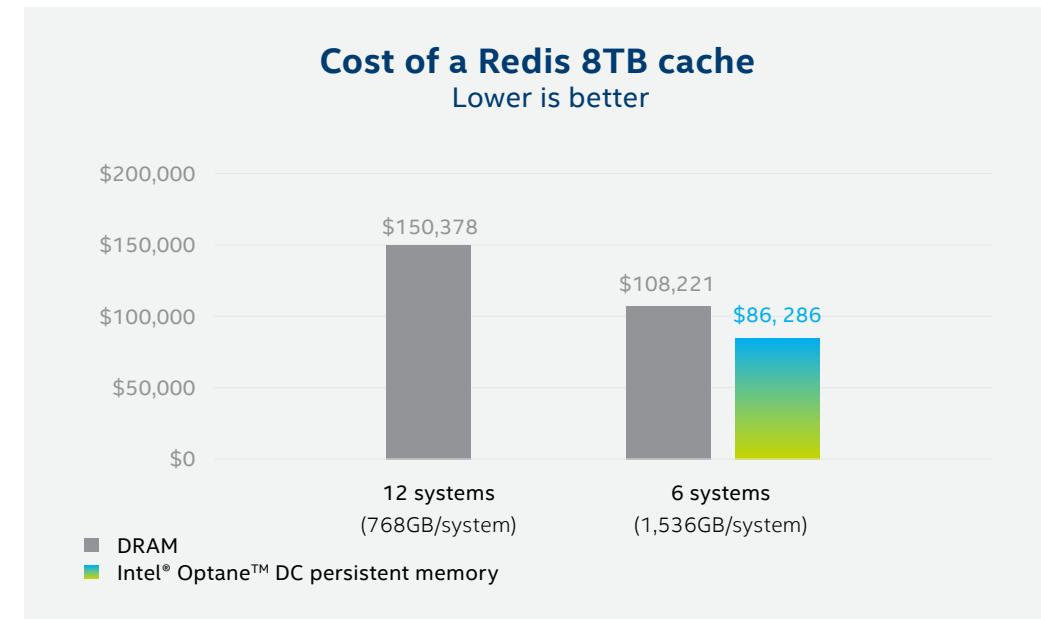


Figure 2: To assemble 8TB cache for Redis requires 12 DRAM systems of 768GB/system, or 6 DRAM or Intel® Optane™ DC persistent memory systems at 1,536GB/system. Compared to either of the DRAM options, Intel Optane DC persistent memory is more cost effective. Intel Optane DC persistent memory enables these system cost savings with more than 3 million operations/second and latency of less than 1ms at the 99th percentile².

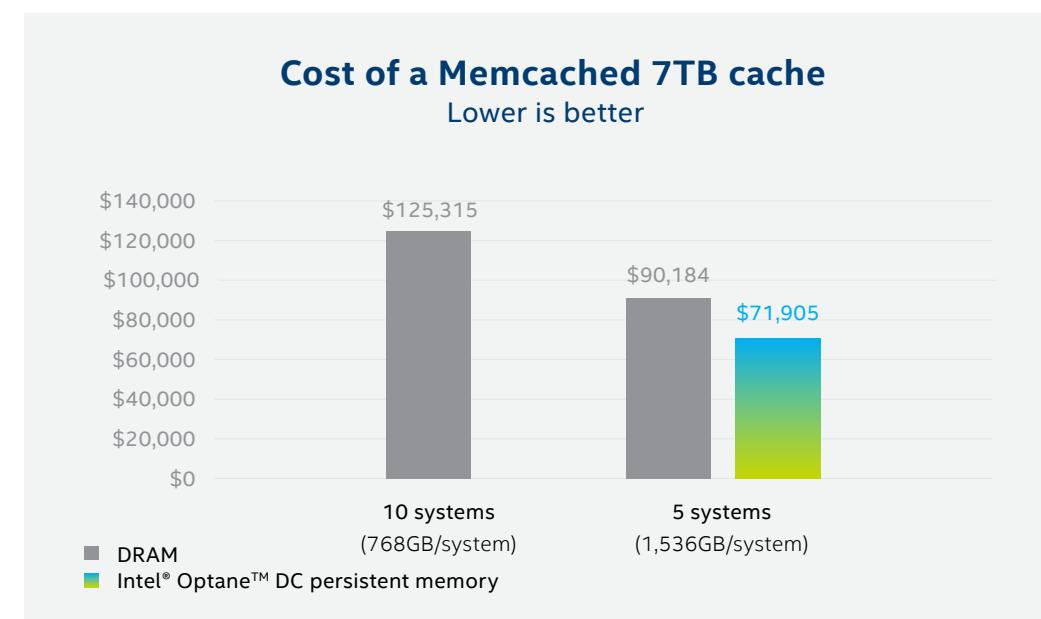


Figure 3: Intel® Optane™ DC persistent memory enables system cost savings with more than 4 million operations/second and latency of less than 1ms at the 99th percentile³.

More responsive storage with App Direct Mode

Intel Optane DC persistent memory can also be used to enhance data storage to increase the speed of persistent reads and writes in the database. Accelerating the main storage can reduce the need for high capacity caching. There are two different ways of deploying Intel Optane DC persistent memory.

App Direct Mode

In App Direct Mode, the database application talks directly to Intel Optane DC persistent memory, mostly bypassing the operating system. This enables the lowest latency possible from persistent memory.

To take advantage of App Direct Mode, the database application you're using needs to be updated for Intel Optane DC persistent memory. Fortunately, many commonly used database applications are being updated.

Some commonly used databases that have been updated for App Direct mode are:

- SAP HANA
- Oracle Exadata
- Aerospike

Storage Over App Direct Mode

You can use Storage Over App Direct Mode to run applications that haven't been updated for Intel Optane DC persistent memory. In this mode, the operating system has been programmed to treat Intel Optane DC persistent memory as a block storage device.

Storage Over App Direct Mode can accelerate performance for databases that have not yet been updated for App Direct Mode, but the speed improvements will not be as great as they would have been under App Direct Mode, and results may vary depending on the database you are using. However, two databases provide illustrative examples (see Figures 4 and 5). Again, no code changes were required.



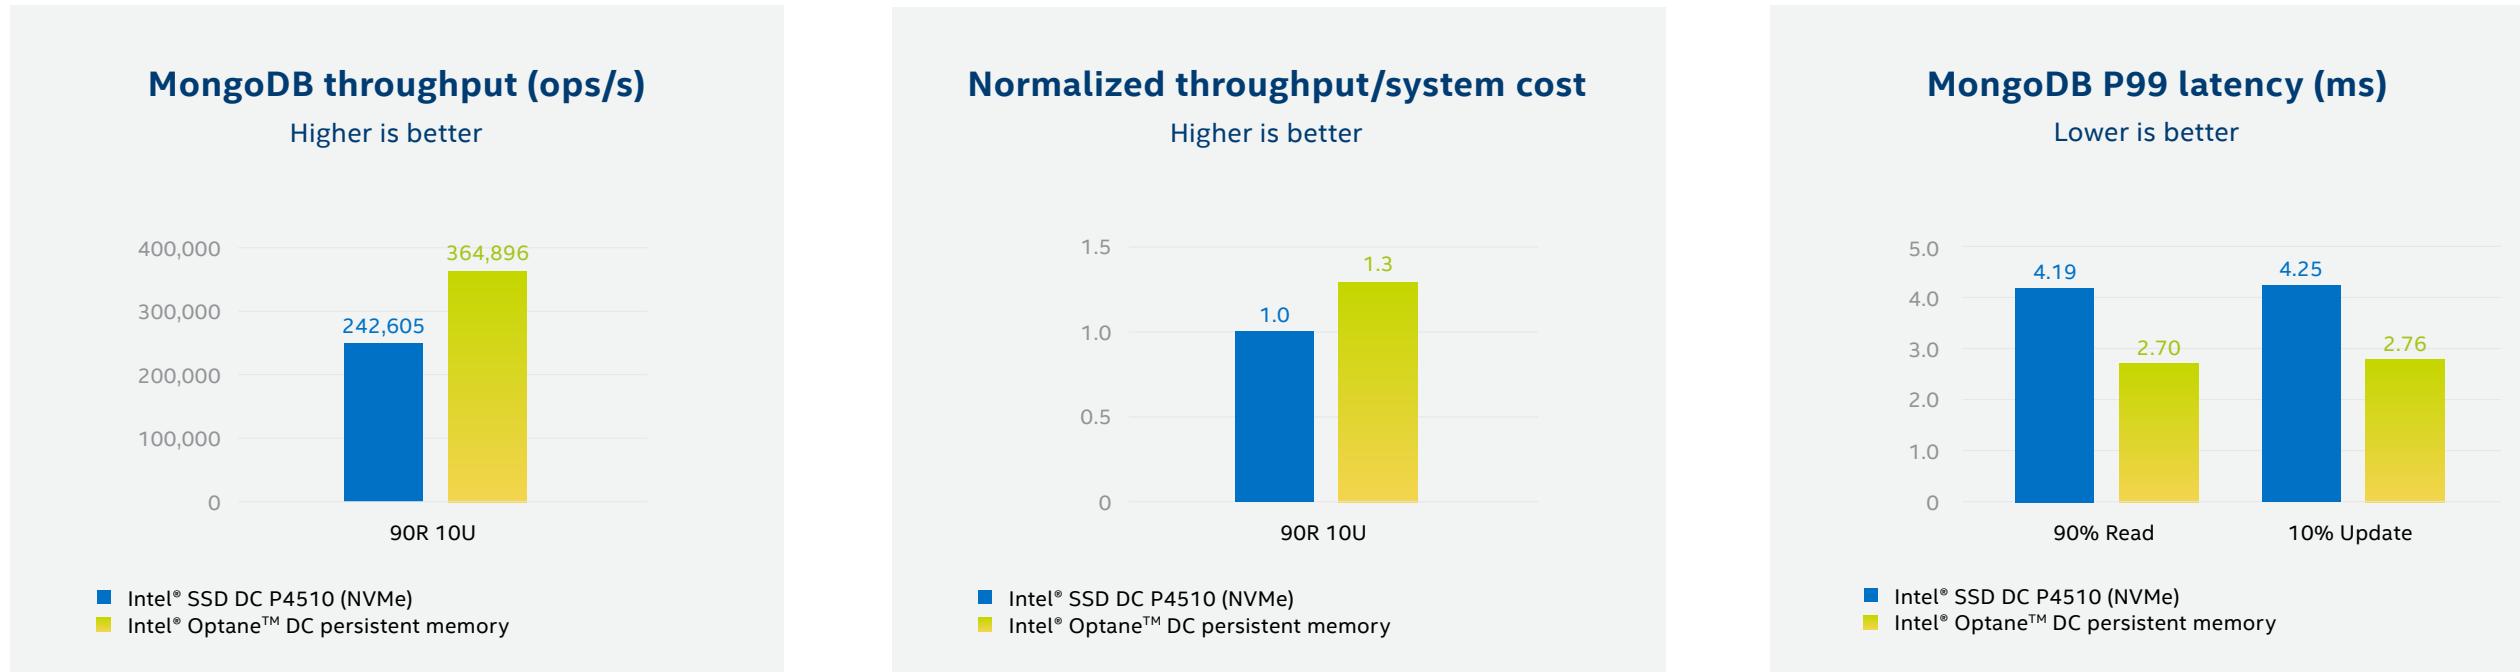


Figure 4: Improving throughput and reducing latency for MongoDB using Intel® Optane™ DC persistent memory in Storage Over App Direct Mode⁴.

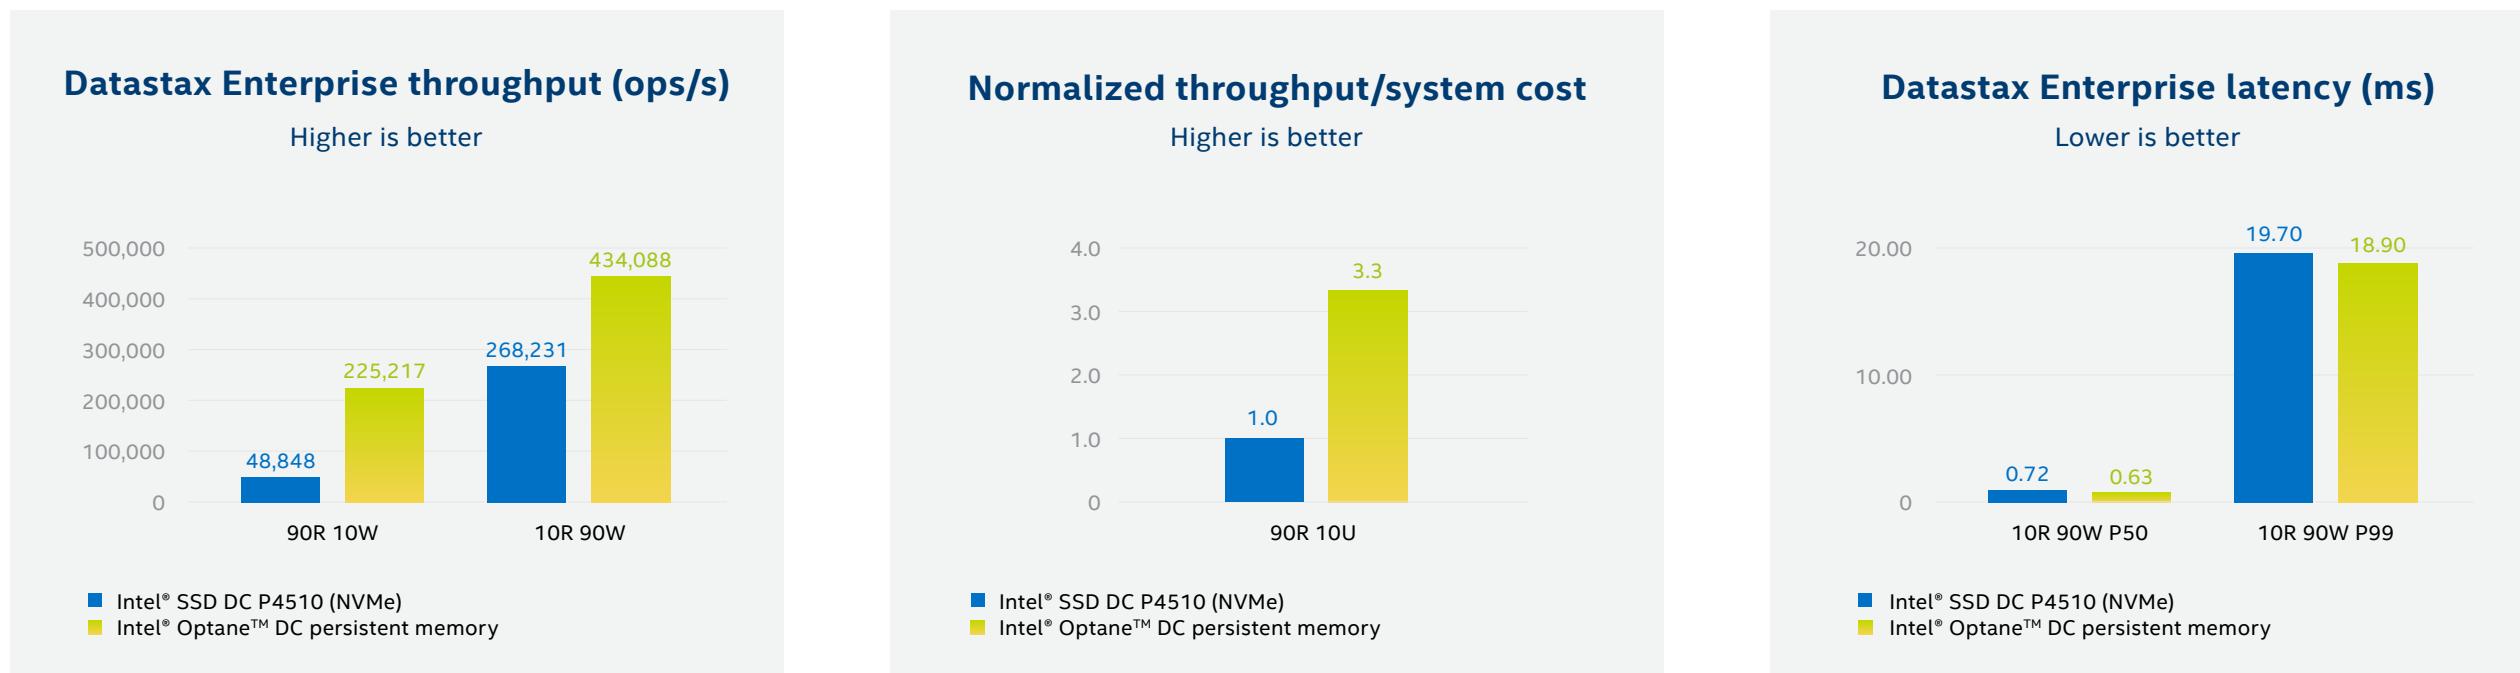


Figure 5: Intel® Optane™ DC persistent memory improves throughput and reduces latency for Datastax Enterprise based on Cassandra⁵.

Crawl, walk, run

You can think of any technology adoption as having the three stages of crawl, walk and run. At first, you'll learn how to get things moving. Once you've mastered the basics (crawl) you can increase your confidence and capabilities (walk). Ultimately, you'll want to achieve the best results the technology is capable of (run).

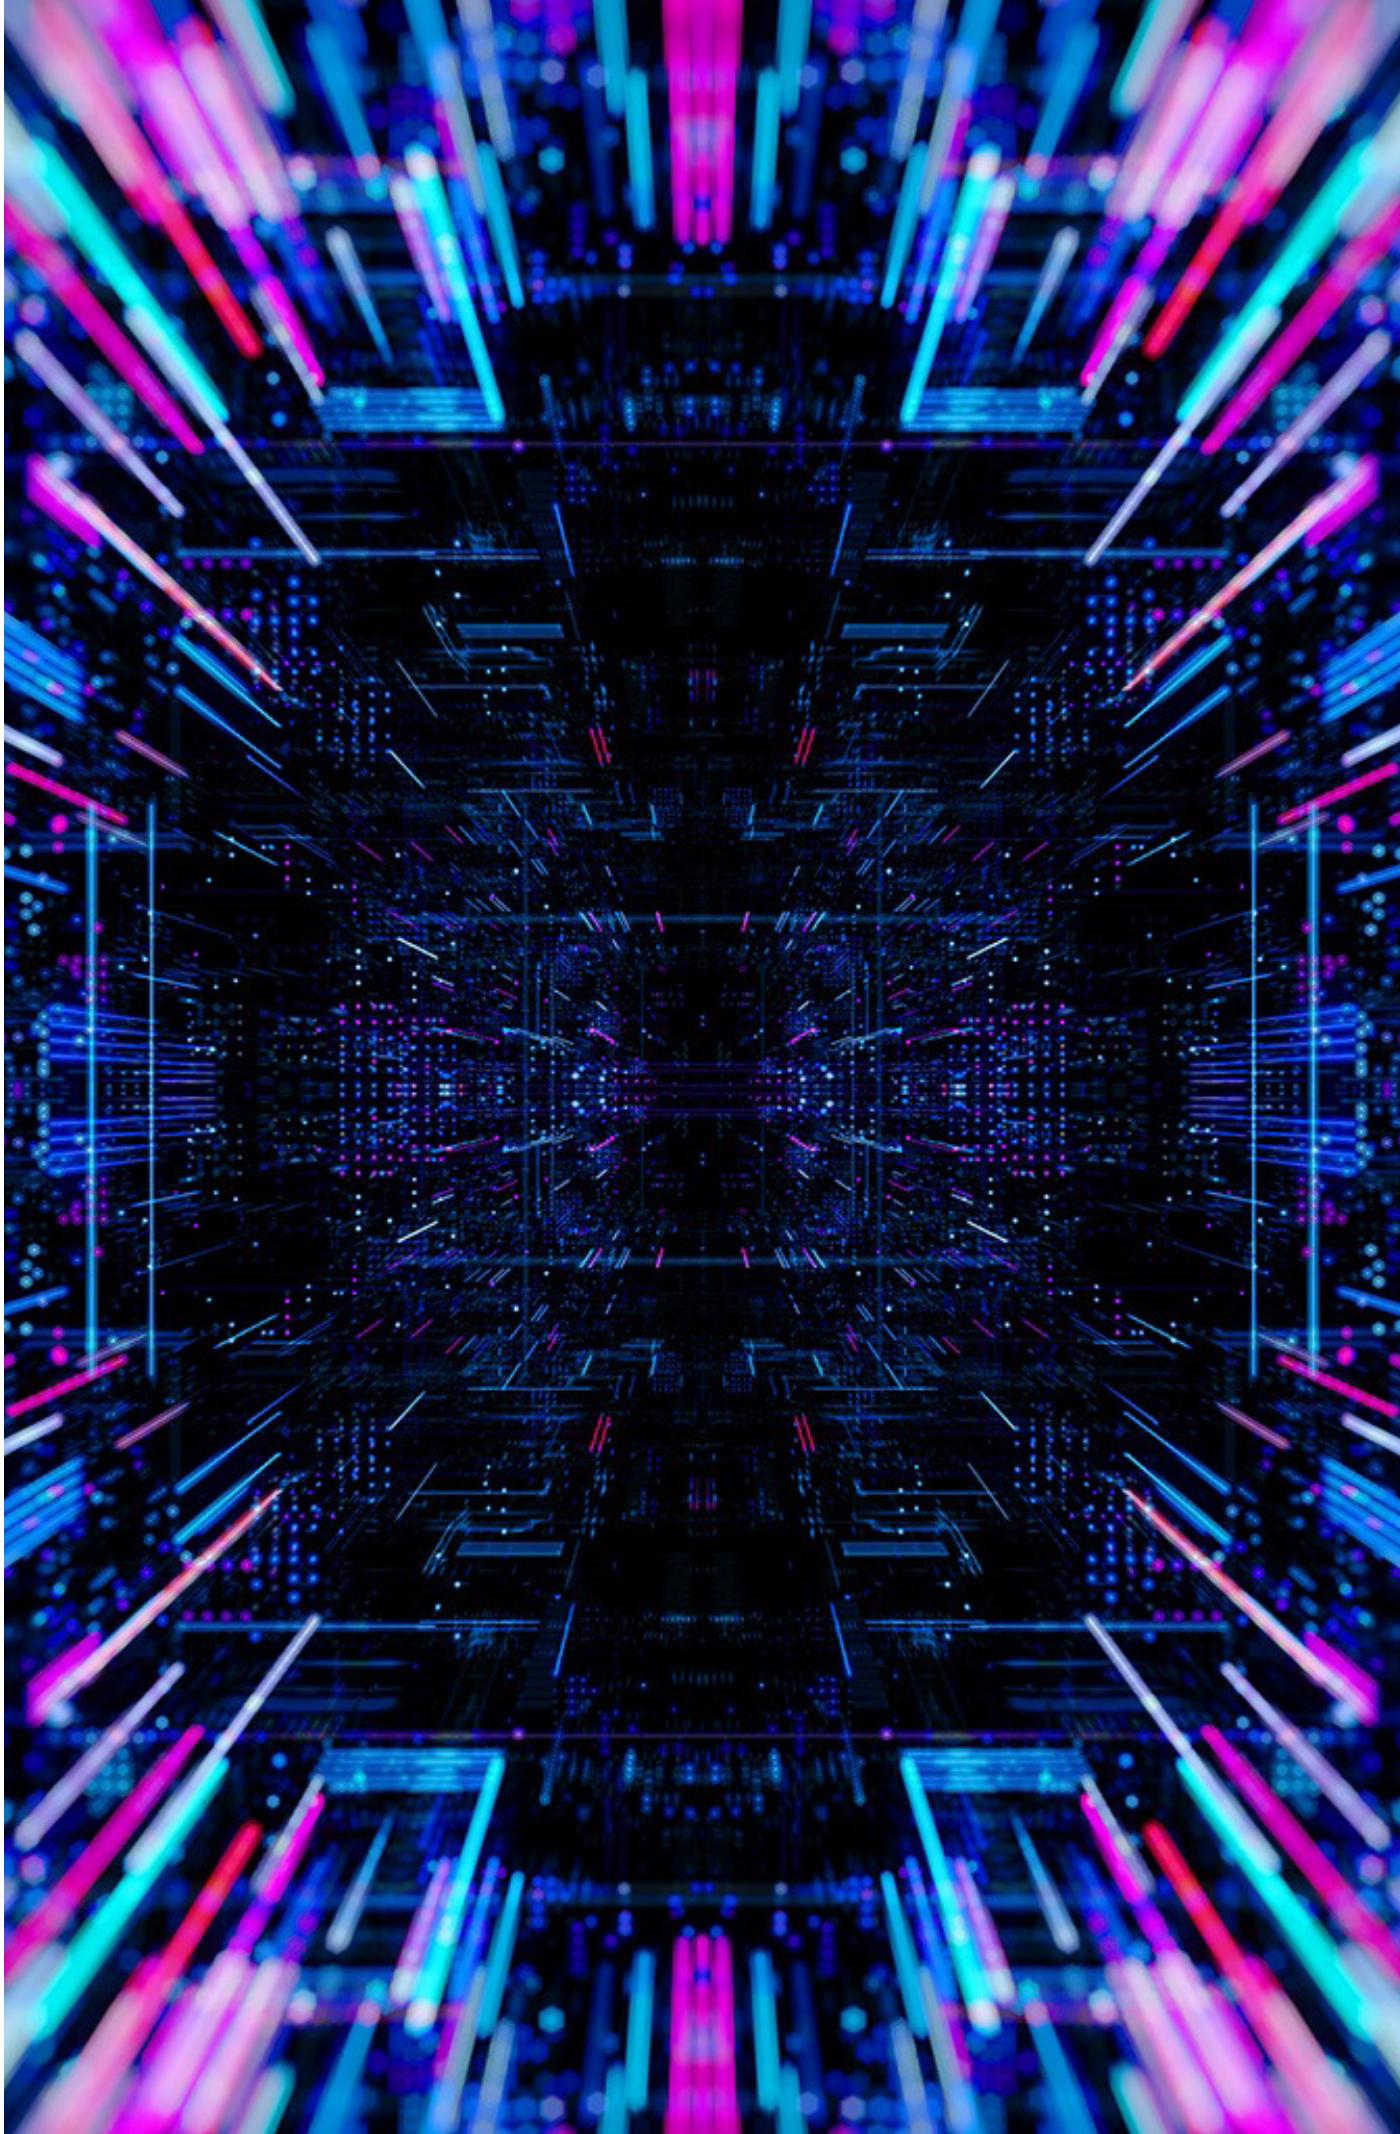
The modes we've looked at so far can help you to start integrating Intel Optane DC persistent memory into your database and enable you to get results relatively quickly. They cover the "crawl" and "walk" stages.

To get to the next level (run), think of how data persistence and large capacities at near-memory speeds can be used to tackle the business problems you are facing. Using App Direct mode, you can develop new, custom-built applications and services tailored to your business users' needs that get optimized performance from Intel Optane DC persistent memory.

The role of SSDs

While Intel Optane DC persistent memory is transformative, there is still a role for SSDs in a database system, due to their high capacity and affordability compared to memory technologies. They are not as fast as Intel Optane DC persistent memory, but they can be used to tier warm data in a way that is more cost effective. They can also be used to provide a solution for databases that are not as performance sensitive and so will not necessarily benefit from the low latency of Intel Optane DC persistent memory.

Compared to hard disks, SSDs are faster and more reliable. The Intel® SSD Data Center Family for NVMe is useful for high-performing databases. The NVMe interface outstrips the performance of SAS/SATA SSDs by optimizing hardware and software to take full advantage of NVMe SSD technology. The Intel® SSD Data Center Family for PCIe provides up to six times faster data transfer speed than 6Gbps SAS/SATA SSDs⁶. Intel's wide range of NVMe SSDs have undergone rigorous quality and compatibility testing. The NVMe interface underpins both Intel Optane DC SSDs and Intel QLC 3D NAND SSDs. Intel Optane DC SSDs offer improved speed and endurance compared to NAND SSDs.



Acceleration technologies

Optimizing database performance isn't all about storage and memory. You can also use acceleration technologies to speed up your database systems.

Deploying FPGAs

Field programmable gate arrays (FPGAs) are semiconductors that can be reprogrammed, creating hardware circuitry customized to support specific processes. They can significantly accelerate workload performance.

The Intel® Programmable Acceleration Card (Intel® PAC) helps simplify deployment in the data center. The first board in the family is based on the Intel® Arria® 10 GX FPGA (Intel® PAC with Intel® Arria® 10 GX FPGA). The 2nd generation board, Intel® FPGA Programmable Acceleration Card (Intel® FPGA PAC) D5005, is based on the Intel® Stratix® 10 SX FPGA and is now available. Devices in the Intel PAC family will be able to share the same accelerator code.

Reprogramming FPGAs has historically taken time and expertise that enterprises might not have access to in-house. Intel's software stack makes the deployment of FPGAs simpler (see Figure 6). The acceleration stack can be used in conjunction with accelerator code from Intel ecosystem players to streamline implementation. This effectively creates off-the-shelf solutions for different workloads.

Here are some solutions tailored for database and analytics applications.

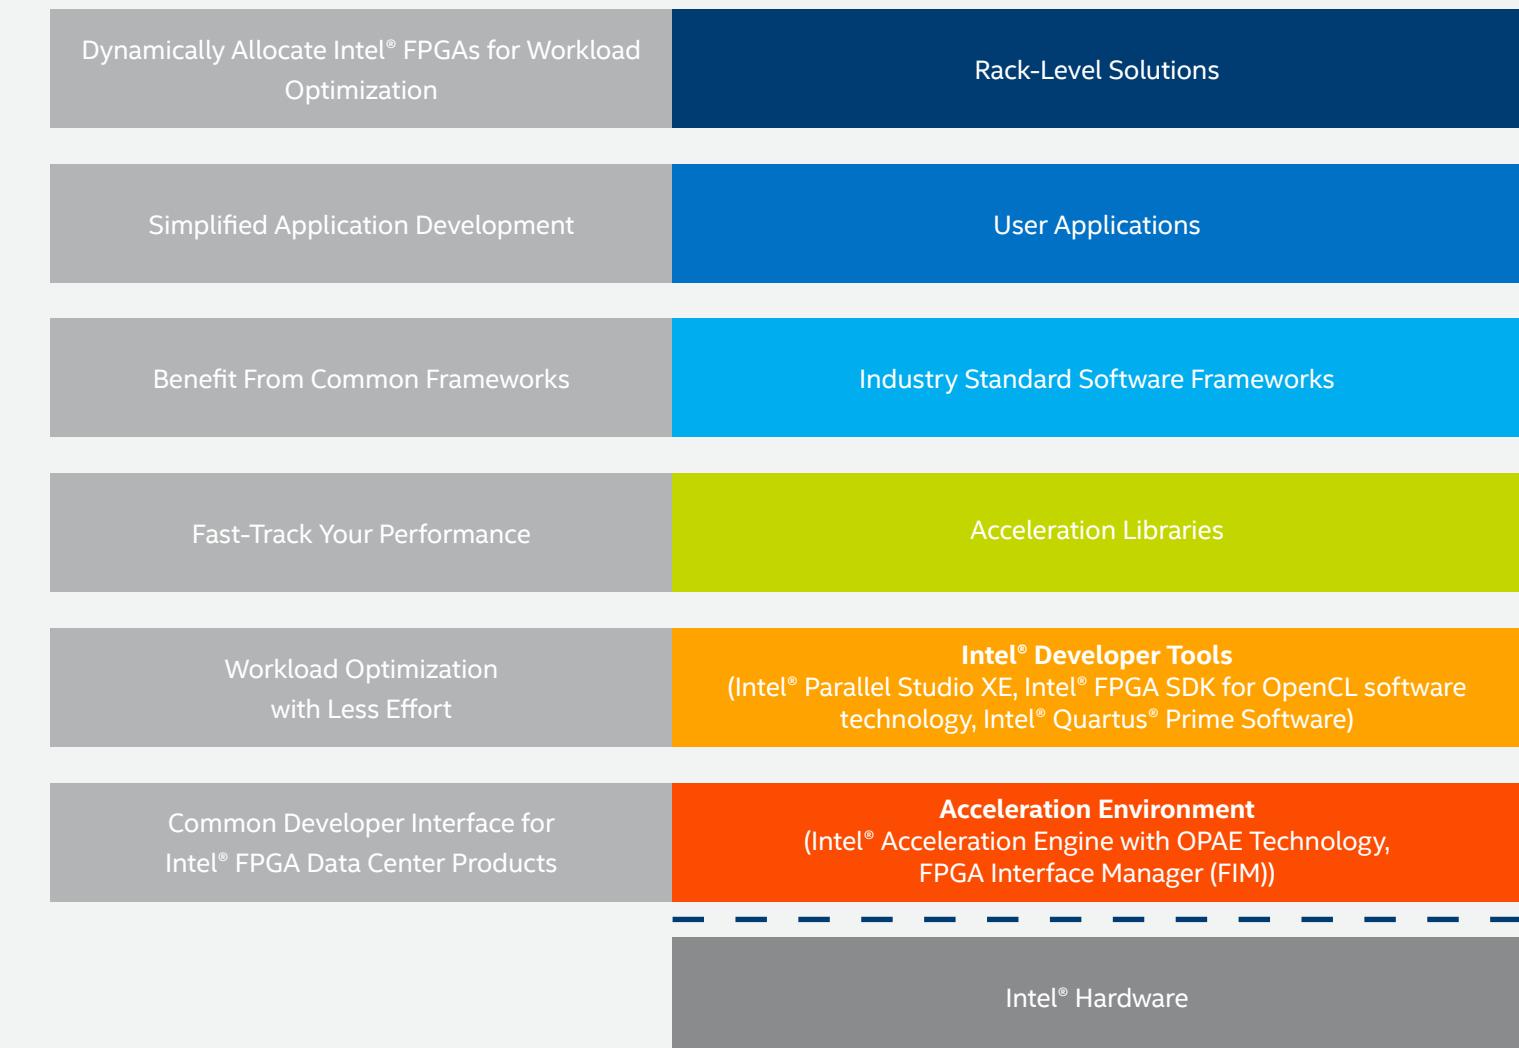


Figure 6: The Intel® Acceleration Stack for FPGAs.

Apache Cassandra

The Data Engine from rENIAC enables FPGAs to be used with standard servers to accelerate data-centric workloads. The Data Proxy, which can accelerate Cassandra's read performance without requiring any changes to application code, has been built on this engine. The solution is hosted on an additional server equipped with an FPGA, which sits between the database clients and the database server.

Apache Spark

Apache Spark is the first platform supported by the Bigstream Hyper-acceleration Layer (HaL). The Apache Spark representation of the computation is translated into a platform-independent dataflow, which HaL can translate into FPGA-optimized code to accelerate database and analytics solutions.

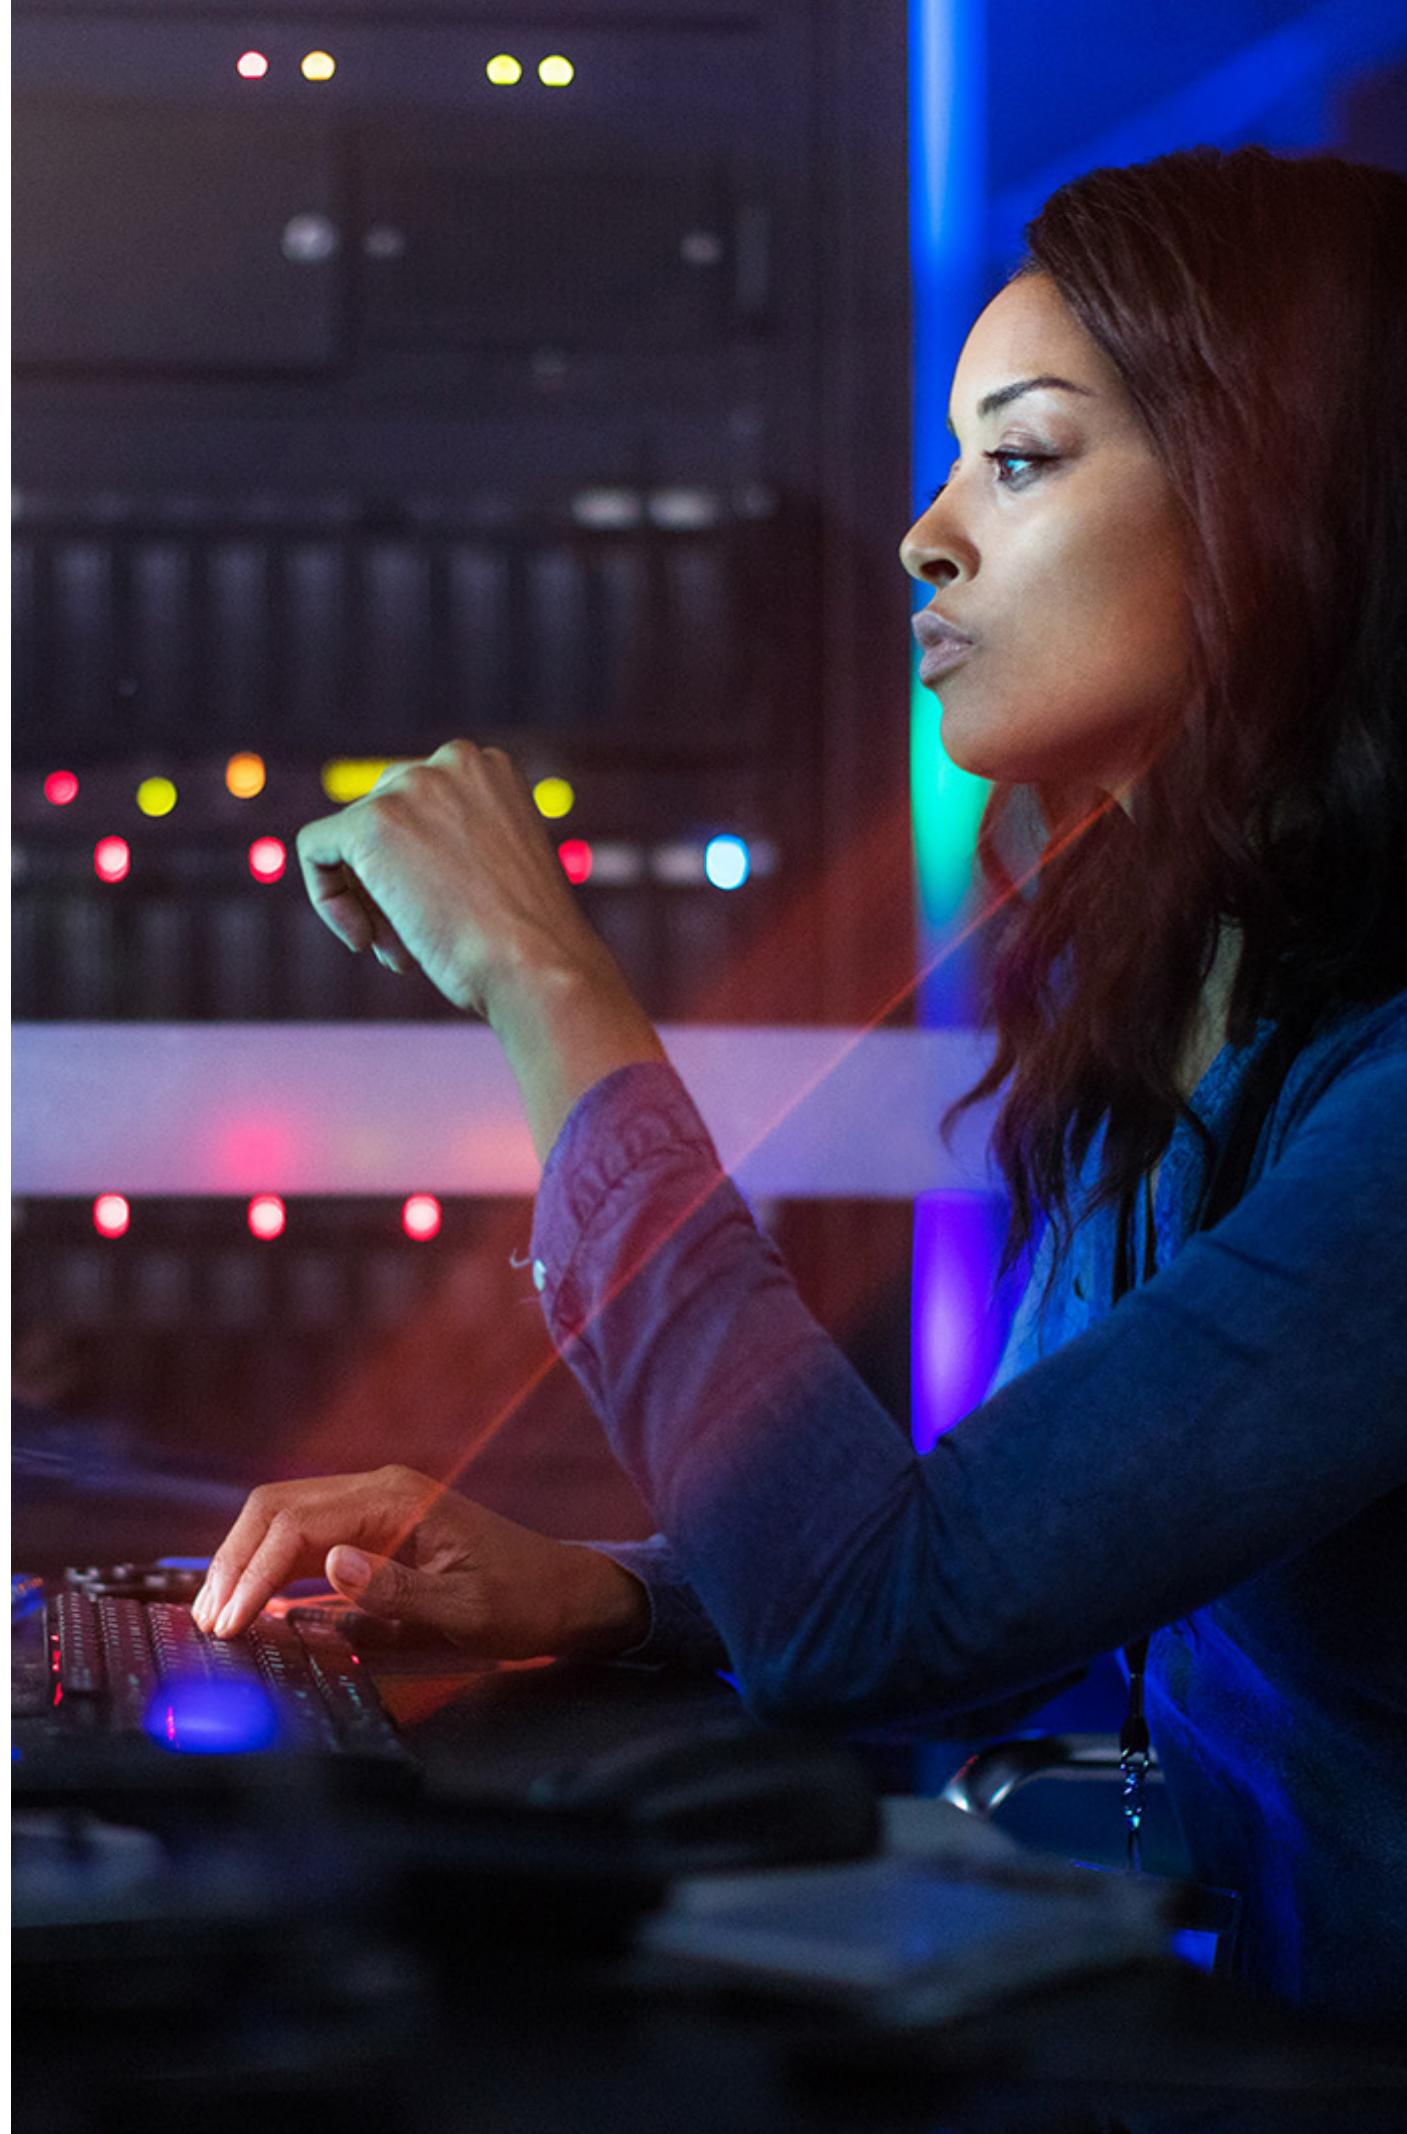
MySQL

Many big data and analytics applications are underpinned by relational databases and SQL. S64DA from Swarm64 uses Intel® FPGAs to accelerate common databases, including PostgreSQL, MariaDB and MySQL. The solution optimizes the data structures of the database and the dataflow, so that each query needs to touch less data, and data that does need to be processed can move through the system more quickly.

Use cases for S64DA include applications that insert data at high velocity (up to millions of rows per second), searching or aggregating ranges in large data sets, and near real-time requirements (sub-second timing between information being received by the database and it being made available to queries).

Intel® Cache Acceleration Technology

The database cache plays an important role in speeding up database accesses, but it's only effective if it's storing the right data. Intel® Cache Acceleration Technology is an affordable way to accelerate the performance of large databases stored on hard disk storage or SSDs. It analyzes data to ensure that it is in the right place, prioritizing the access to hot data and placing colder data on tiered storage. Using this technology with scale-out databases can cut latency and reduce restart time. The technology has recently become open source as the Open CAS Framework.



Conclusion

Boosting connectivity

After swapping out slower storage for persistent memory, networking often becomes the new bottleneck. To achieve the full potential the server is capable of, it is important that the server can send and receive data without adding latency into the processing cycle. Data centers reliant on 1GbE networking may find their server performance is limited by the network.

To solve this problem, you can upgrade networking hardware to improve performance.

Intel has a portfolio of leading networking technologies that could be useful, including the Intel® Ethernet Adaptor XXV710-SA2, which delivers excellent performance for 25GbE connectivity. It features Intel® Virtualization Technology (Intel® VT), which improves I/O performance in virtualized server environments, and has backwards compatibility at 10GbE with the cabling your data center might have been using for 1GbE, while providing a smooth upgrade path to 25GbE.

Intel® Ethernet Flow Director (Intel® Ethernet FD) is a built-in traffic steering capability which directs traffic into specific queues to avoid context switching within the CPU. It can significantly increase the number of transactions per second and reduce latency for applications such as Memcached.

When it comes to deployment, the Intel® Ethernet Network Adaptor XXV710 implements a built in Link Establishment State Machine (LESM) which allows it to support all 25GbE link types.

Intel Optane DC persistent memory can be used to improve memory and storage in the data center, and a variety of acceleration technologies are available for accelerating database performance including FPGAs and Intel Cache Acceleration Technology. By providing a high-performance foundation for their databases, IT leaders can ensure their applications deliver the results their users need, quickly.

Discover more resources:
visit intel.com/optane

Learn more

- [Intel® Optane™ DC Persistent Memory](#)
- [2nd Generation Intel® Xeon® Scalable Processor](#)
- [Intel® Programmable Acceleration Card with Intel Arria® 10 GX FPGA](#)
- [Intel® FPGA Programmable Acceleration Card D5005](#)
- [Intel® FPGA Acceleration Hub](#)
- [Intel® Cache Acceleration Technology](#)
- [Intel® Ethernet Network Adapters](#)
- [Intel® Solid State Drives](#)

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 complete information visit <https://www.intel.com/benchmarks>.

¹ <https://www.intel.com/content/www/uk/en/architecture-and-technology/avx-512-overview.html>

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

Intel® AVX 2.0 delivers 16 double precision and 32 single precision floating point operations per second per clock cycle within the 256-bit vectors, with up to two 256-bit fused-multiply add (FMA) units.

² Running Redis 4.0.13 as a Clustered LRU Cache with 90%/10% Read/Write profile on IO intensive queries, and provided to you for informational purposes. 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.

Performance results are based on testing as of April 14, 2019 Intel® Optane™ DC Persistent Memory pricing & DRAM pricing as of March 13, 2019.

Common system configuration details for all Redis tests: single node, 2x Intel® Xeon® Gold 6248 Processor, 20 cores, HT On, Turbo ON, 1x S3520 SATA 240GB boot drive, 1x dual port XXV710-DA2 25Gb/s NIC, BIOS: SE5C620.86B.0D.01. 0374.013120191836 (ucode:0x400001C), OS: Ubuntu 18.04 LTS, kernel: Linux 4.15.0-47-generic x86_64, gcc 7.3.0 compiler

Common Testing details for all Redis tests: Redis 4.0.13 deployed bare metal in clustered, sharded LRU cache, 1k byte Value size, numactl used to bind each shard to a cpu socket and memory, evenly distributed across the 2 cpu sockets and across the 2 IP addresses, load generated by Memtier_Benchmark 1.2.15, 90R/10W, Random Distribution, Prepopulate Values with random data, Increase ops/s incrementally until P99 latency exceeds 1ms, record throughput at last value where P99 < 1ms

Config 1: 768GB DRAM: Test by Intel as of 4/14/2019. Total Memory: 768GB DDR4 (24 slots/ 32GB/ 2666 MHz), system consists of 36 shards of 20GB (720GB of cache) score: 3.93M ops/sec with P99 Latency of 0.9ms - create 8TB multi-node cache with 12 server nodes that deliver > 3Mops/sec with P99 Latency < 1ms per system

Config 2: 1,536GB DRAM: Test by Intel as of 4/14/2019. Total Memory: 1.536GB DDR4 (24 slots/ 64GB/ 2666 MHz), system consists of 36 shards of 40GB (1,440GB of cache) score: 3.94M ops/sec with P99 Latency of 0.9ms - create 8TB multi-node cache with 6 server nodes that deliver > 3Mops/sec with P99 Latency < 1ms per system

Config 3: 1,536GB DCPMM: Test by Intel as of 4/14/2019. Total Memory: 192GB DDR4 (12 slots/ 16GB/ 2666 MHz) + Optane DCPMM 1.5TB (12 slots/ 128GB/ 2666MHz) in Memory mode, system consists of 32 shards of 42.2GB (1,350GB of cache) score: 3.24M ops/sec with P99 Latency of 0.9ms - create 8TB multi-node cache with 6 server nodes that deliver > 3Mops/sec with P99 Latency < 1ms per system

³ Results based on tests conducted on pre-production systems running Memcached 1.5.12 instances with 90%/10% Read/Write profile on IO intensive queries, and provided to you for informational purposes. Pricing guidance is subject to change and may revise up or down based on market dynamics. Please contact your OEM/distributor for actual pricing.

Cost of a 7TB cache based on estimated system cost for a given configuration (see configuration details slide) times the number of systems of that configuration required to add up to the target cache size. Performance results are based on testing as of April 14, 2019 and may not reflect all publicly available security updates. No product or component can be absolutely secure.

Intel® Optane™ DC Persistent Memory pricing & DRAM pricing as of March 13, 2019. Pricing referenced in TCO calculations is provided for guidance and planning purposes only and does not constitute a final offer.

Common system configuration details for all Memcached tests: single node, 2x Intel® Xeon® Gold 6248 Processor, 20 cores, HT ON, Turbo ON, 1x S3520 SATA 240GB boot drive, 1x dual port XXV710-DA2 25Gb/s NIC, BIOS: SE5C620.86B.0D.01. 0374.013120191836 (ucode:0x400001C), OS: Ubuntu 18.04 LTS, kernel: Linux 4.15.0-47-generic x86_64, gcc 7.3.0 compiler

Common Testing details for all Memcached tests: Memcached 1.5.12 deployed bare metal instances, 1k byte Value Size, numactl used to bind each instance to a cpu core and memory, evenly distributed across the 2 cpu sockets and across the 2 IP addresses, 80 network IRQs spread across cores, load generated by RPC-perf 3.0.0-pre, 90R/10W, Random Distribution, Prepopulate Values with random data, Increase ops/s incrementally until P99 latency exceeds 1ms, record throughput at last value where P99 < 1ms

Config 1: 768GB DRAM: Test by Intel as of 4/14/2019. Total Memory: 768GB DDR4 (24 slots/ 32GB/ 2666 MHz), system consists of 60 instances to create 700GB of cache. score: 4.43M ops/sec with P99 Latency of 0.9ms - create 7TB multi-node cache with 10 server nodes with > 4Mops/sec with P99 Latency < 1ms per system

Config 2: 1,536GB DRAM: Test by Intel as of 4/14/2019. Total Memory: 1.536GB DDR4 (24 slots/ 64GB/ 2666 MHz), system consists of 60 instances to create 1,400GB of cache. score: 4.48M ops/sec with P99 Latency of 0.6ms - create 7TB multi-node cache with 5 server nodes with > 4Mops/sec with P99 Latency < 1ms per system

Config 3: 1,536GB DPCMM: Test by Intel as of 4/14/2019. Total Memory: 192GB DDR4 (12 slots/ 16GB/ 2666 MHz) + Optane DCPMM 1.5TB (12 slots/ 128GB/ 2666MHz) in Memory mode, system consists of 60 instances to create 1,400GB of cache. score: 4.19M ops/sec with P99 Latency of 0.9ms - create 7TB multi-node cache with 5 server nodes that deliver > 4Mops/sec with P99 Latency < 1ms per system

⁴ Performance results are based on testing as of May 1, 2019 and may not reflect all publicly available security updates. No product or component can be absolutely secure.

Results based on tests conducted on pre-production systems running MongoDB 4.0.4 as a sharded database with 90%/10% Read/Write profile on IO intensive queries, and provided to you for informational purposes.

Intel® Optane™ DC Persistent Memory pricing & DRAM pricing as of March 13, 2019. 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.

Normalized throughput/system cost is the throughput of a given configuration divided by system cost – normalized to the baseline NVME system

Common system configuration details for all MongoDB tests: single node, 2x Intel® Xeon® Gold 6248 Processor, 20 cores, HT On, Turbo ON, 1x S3520 SATA 240GB boot drive, 1x dual port XXV710-DA2 25Gb/s NIC, BIOS: SE5C620.86B.0D.01.0416.030220191752 (ucode:0x4000021), OS: Ubuntu 18.04 LTS, kernel: Linux 4.15.0-48-generic x86_64, gcc 7.4.0 compiler

Common Testing details for all MongoDB tests: MongoDB Community Edition 4.0.4 deployed bare metal in sharded database, create 4 shards of 350GB with 50million operations on each shard, 1k row size, numactl used to bind each shard to a cpu socket, evenly distributed across the 2 cpu sockets , load generated by YCSB-0.15.0 (PerfKitBenchmark), 90R/10W, Uniform Distribution.

Config 1: NVME: Test by Intel as of 5/1/2019. Total Memory: 384GB DDR4 (12 slots/ 32GB/ 2666 MHz), system 4x P4510 NVME 2TB configured in RAID 0, score: 101,948 ops/sec with P99 Latencies of 7.36ms (read) and 21.15ms (update) - Config 2: DPCMM: Test by Intel as of 5/1/2019. Total Memory: 384GB DDR4 (12 slots/ 32GB/ 2666 MHz), + Optane DCPMM 1.5TB (12 slots/ 128GB/ 2666MHz) in Storage over App Direct mode configured using DAX, score: 322,103 ops/sec with P99 Latencies of 2.72ms (read) and 5.51ms (update) – 2.3 times savings in throughput/per system cost compared to NVME configuration

⁵ Performance results are based on testing as of April 14, 2019 by DataStax and may not reflect all publicly available security updates. No product or component can be absolutely secure.

Results based on tests conducted on pre-production systems running DataStax Enterprise 6.7.2 and DataStax Distribution of Apache Cassandra 5.1.13 as a 4 node clustered database with 90%/10% and 10%/90% Read/Write profile on IO intensive queries, and provided to you for informational purposes. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Intel® Optane™ DC Persistent Memory pricing & DRAM pricing as of March 13, 2019. 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.

Normalized throughput/system cost is the throughput of a given configuration (see configuration slide for details) divided by system cost – normalized to the baseline cost of the NVME configuration

Common system configuration details for all Cassandra tests: 4 node cluster, 2x Intel® Xeon® Gold 6248 Processor, 20 cores, HT On, Turbo ON, 1x S3520 SATA 240GB boot drive, 1x dual port XXV710-DA2 25Gb/s NIC, BIOS: SE5C620.86B.0D.01.0416.030220191752 (ucode:0x4000021), OS: Ubuntu 18.04 LTS, kernel: Linux 4.15.0-48-generic x86_64, gcc 7.4.0 compiler

Common Testing details for all Cassandra tests: DataStax Enterprise 6.7.2 deployed bare metal in 4 node clustered database, create 650GB database from snapshot and ramp threads til cpu is 85-90% utilized. Load generated by EBDSE IoT workload profile, 90R/10W, 50R/50W, 10R 90W workload compositions. Uniform Distribution.

Config 1: NVME: Test by DataStax as of 5/1/2019. Total Memory: 384GB DDR4 (12 slots/ 32GB/ 2666 MHz), system 3x P4510 NVME 2TB configured in RAID 0, score: 101,948 ops/sec with P99 Latencies of 7.36ms (read) and 21.15ms (update) - Config 2: DPCMM: Test by DataStax as of 5/1/2019. Total Memory: 384GB DDR4 (12 slots/ 32GB/ 2666 MHz), + Optane DCPMM 1.5TB (12 slots/ 128GB/ 2666MHz) in Storage over App Direct mode configured using DAX, score: 322,103 ops/sec with P99 Latencies of 2.72ms (read) and 5.51ms (update) – 2.3 times savings in throughput/per system cost compared to NVME configuration

⁶ Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Configurations: Performance claims obtained from data sheet, sequential read/write at 128k block size for NVMe* and SATA, 64k for SAS. Intel® SSD DC P3700 Series 2TB, SAS Ultrastar SSD1600MM, Intel® SSD DC S3700 Series SATA 6Gbps. Intel® Core™ i7-3770K CPU @ 3.50GHz, 8GB of system memory, Windows* Server 2012, IOMeter. Random performance is collected with 4 workers each with 32 QD. <https://www.intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/intel-ssd-dc-family-for-pcie.html>

Performance results are based on testing as of the date set forth in the configurations and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

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.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Optimization Notice: Intel's 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.

Notice Revision #20110804

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer, or learn more at intel.com.

Intel, the Intel logo, and other Intel Marks are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

Other names and brands may be claimed as the property of others.

