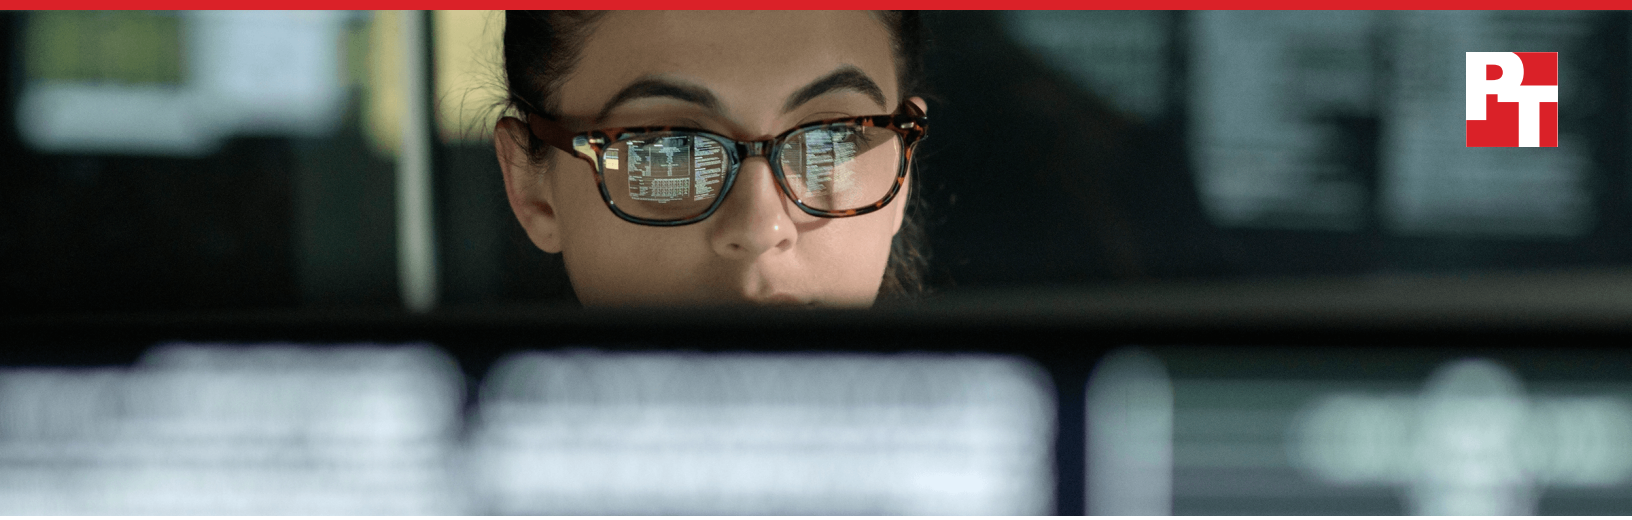

Improve Aerospike Database performance and predictability by leveraging Intel® Ethernet 800 Series Network Adapters with Application Device Queues (ADQ)

In Dell EMC PowerEdge R740xd servers, enabling ADQ boosted Aerospike performance and led to near-linear scaling in both DRAM and Intel® Optane™ persistent memory configurations

Scale-out applications that rely on communication between servers depend on fast networking performance. Now, a feature of Intel® Ethernet 800 Series Network Adapters, Application Device Queues (ADQ), offers a streamlined path for data access.

In the Principled Technologies data center, we set up three Dell EMC™ PowerEdge™ R740xd server nodes and used a built-in benchmark tool to assess Aerospike Database 5 (a NoSQL database platform) performance with ADQ enabled and ADQ disabled at multiple node counts. Our tests showed that enabling ADQ with the Intel Ethernet 800 Series Network Adapter provided a significant improvement in throughput, latency, and predictability for Aerospike databases—which helps deliver responsive experiences for end users. Enabling ADQ also offered near-linear scaling for Aerospike Database applications as we added server nodes.

To make the most of your Aerospike Database workloads on Dell EMC PowerEdge R740xd servers, enabling Application Device Queues on Intel Ethernet 800 Series Network Adapters can give these high-priority database applications a boost.

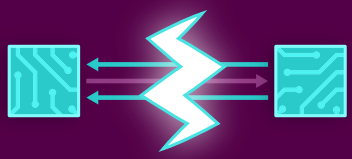

Increase Aerospike Database throughput by up to 64%

Near-linear scaling for Aerospike databases

Increase predictability of response times by up to 41% on Aerospike databases

The importance of being predictably fast

Distributed databases—where data resides on more than one node, typically many physical nodes—rely on more than just server processing power to ensure fast performance; the network connecting this data must be fast. Database performance must be fast, but also have predictable response times, limiting slow outlier requests that add up to latency problems as the applications scale. Dell EMC PowerEdge R740xd servers utilize Intel Ethernet 800 Series Network Adapters to meet these twin challenges.

Intel Ethernet 800 Series Network Adapters offer Application Device Queues, which route data to dedicated hardware queues to decrease latency, improve response time predictability, and ultimately deliver higher throughput for your critical applications. By creating what are essentially express lanes for high-priority data, enabling ADQ on the Intel Ethernet 800 Series can help you more easily meet service-level agreements (SLAs) and scale applications confidently while providing very low latencies.

To show how ADQ can benefit applications, we created a testbed with three Dell EMC PowerEdge R740xd servers and compared Aerospike Database performance with ADQ enabled and with ADQ turned off (see Figure 1). (To see other applications that can utilize ADQ, visit www.intel.com/adq.) We tested the Aerospike Database application using a benchmark built into the Aerospike platform: the Aerospike C benchmark. This simple benchmark allows users to insert records into the database and run tests with a variety of read/write mixes (we used 100% reads in our tests) and thread counts to assess overall Aerospike Database performance. The benchmark reports its results in transactions per second (TPS). For complete testing details and results, read the science behind the report.

About Aerospike Database 5

In our tests, we used Aerospike Database Enterprise Edition build 5.3.0.2. Aerospike Database 5 is a NoSQL database platform that supports multiple deployment options on premises or in the cloud. According to Aerospike, "[the database platform] supports two distinct active-active deployment models, multi-site clustering (MSC) and cross datacenter replication (XDR) for building a global data hub."[1] To learn more about Aerospike Database, visit https://www.aerospike.com/products/database-platform/.

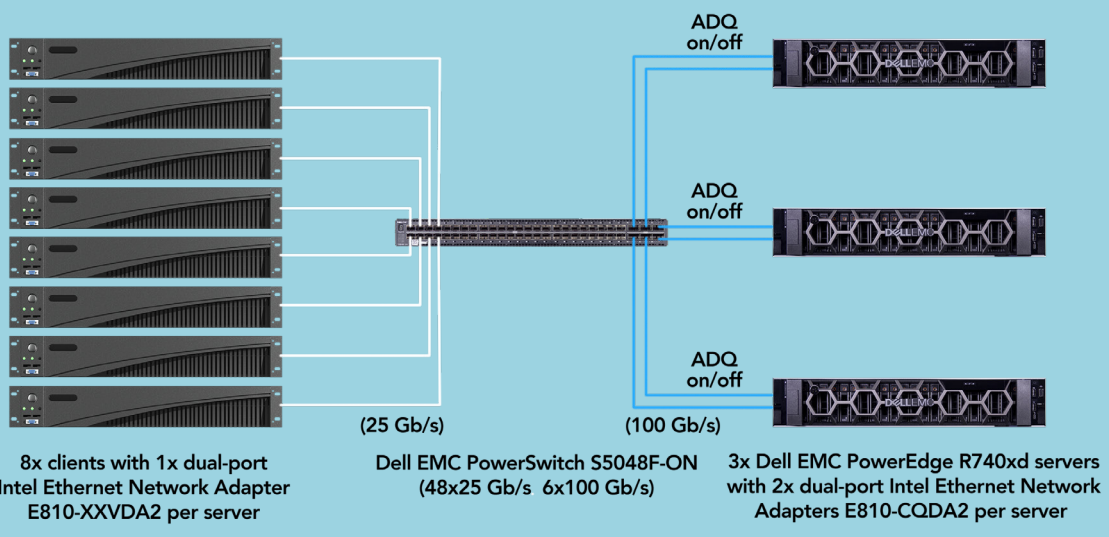

Figure 1: Our testbed setup for Aerospike Database testing. Source: Principled Technologies.

Throughput and scaling

Throughput is a measure of how much work can be done in a certain time—how many orders, operations, Megabytes, and so on. The more application throughput a server can handle with good response times, the more customer requests it can process before you need to support additional nodes. In tests of Aerospike databases, we compared the application throughput with ADQ enabled vs. ADQ disabled to show how using ADQ can prevent network bottlenecks and enable more customer requests to get through. Both ADQ enabled and disabled configurations use optimization via NUMA pinning.

As Figure 2 shows, enabling ADQ dramatically improved the number of random read transactions per second that the Dell EMC PowerEdge R740xd could handle at multiple server node counts, increasing read throughput by up to 64.8 percent.

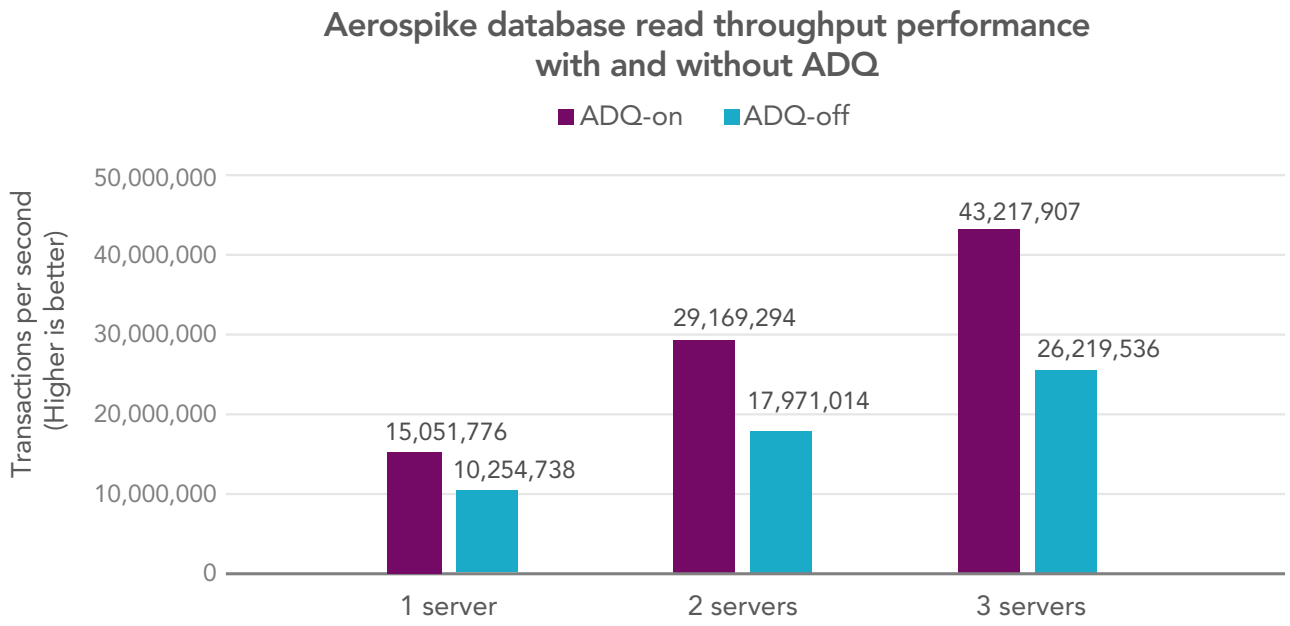

Figure 2: Aerospike database read throughput performance, in transactions per second, of the Dell EMC PowerEdge R740xd at multiple server node counts with Intel Ethernet 800 Series Network Adapters with ADQ on and ADQ off. Source: Principled Technologies.

About 2nd Generation Intel Xeon Scalable processors

The 2nd Generation Intel Xeon® Scalable processor platform features a wide range of processor types to support the workloads you run, including Bronze, Silver, Gold, and Platinum. According to Intel, the 2nd Generation Intel Xeon Scalable platform can handle a variety of workloads, including enterprise, cloud, HPC, storage, and communications. The new processor line also supports a new memory and storage technology to further accelerate workloads, Intel Optane DC persistent memory. To learn more about the Intel Xeon Scalable processor family, visit https://www.intel.com/content/www/us/en/products/docs/processors/xeon/2nd-gen-xeon-scalable-processors-brief.html.

Ideally, an application would scale linearly as identical servers are added—meaning that two servers do exactly twice the work of a single server, three servers do exactly three times the work of a single server, and so on. Because servers communicating with one another requires some amount of overhead to account for their ability to work in concert, perfect linear scaling is generally not possible. Choosing hardware and settings that closely approach this ideal ensure that you make the most of your available resources. As Figure 3 shows, enabling ADQ did more than just provide better throughput overall—enabling ADQ also improved scaling noticeably as we added identical server nodes compared to the same servers with ADQ disabled. At three nodes, the ADQ-enabled configuration deviated just 4.3 percent from linear scaling.

Using ADQ in conjunction with the Intel Ethernet 800 Series Network Adapters to route traffic to dedicated hardware queues allowed the Dell EMC PowerEdge R740xd servers to spend more resources on ensuring high throughput for the Aerospike Database application, which could, at large scale, ultimately lead to the purchase and maintenance of fewer servers to support user requests.

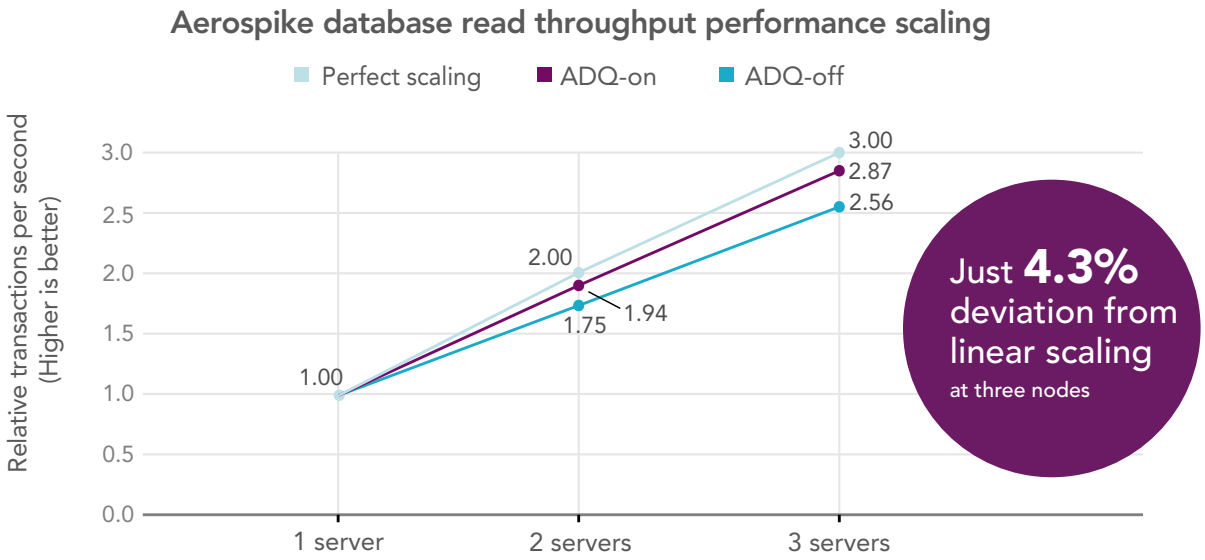

Figure 3: Scaling for Aerospike databases over multiple Dell EMC PowerEdge R740xd server nodes with Intel Ethernet 800 Series Network Adapters with ADQ on and ADQ off. We included a line representing perfect scaling to show how close each configuration came. Source: Principled Technologies.

Latency and predictability

Even the best performing servers with high throughput and low average latency experience outlier requests that fall outside of acceptable response times. At low server counts, these slowest response times, known as tail latency, are not likely to significantly affect user experience. However, in deployments at scale—with many servers supporting the same application—these slow responses can reduce the ability of your organization to meet SLAs and adequately support users. Reducing jitter, the term for the variation in latency, can help you predictably scale Aerospike Database applications to ensure consistent performance. In our tests, enabling ADQ in addition to NUMA pinning, which assigns memory and CPU usage of Aerospike databases to a single NUMA node, helped to reduce jitter and ensure predictable, consistent performance.

Figure 4 compares the response times for 99th percentile (P99) latency—which means the response time that 99 percent of all requests achieved. Lower response times can help enable better overall performance, and having a lower P99 latency means that nearly all response times from all requests were lower, which improves predictability at scale. With a single server node, turning on ADQ resulted in a 41.7 percent improvement in latency predictability.

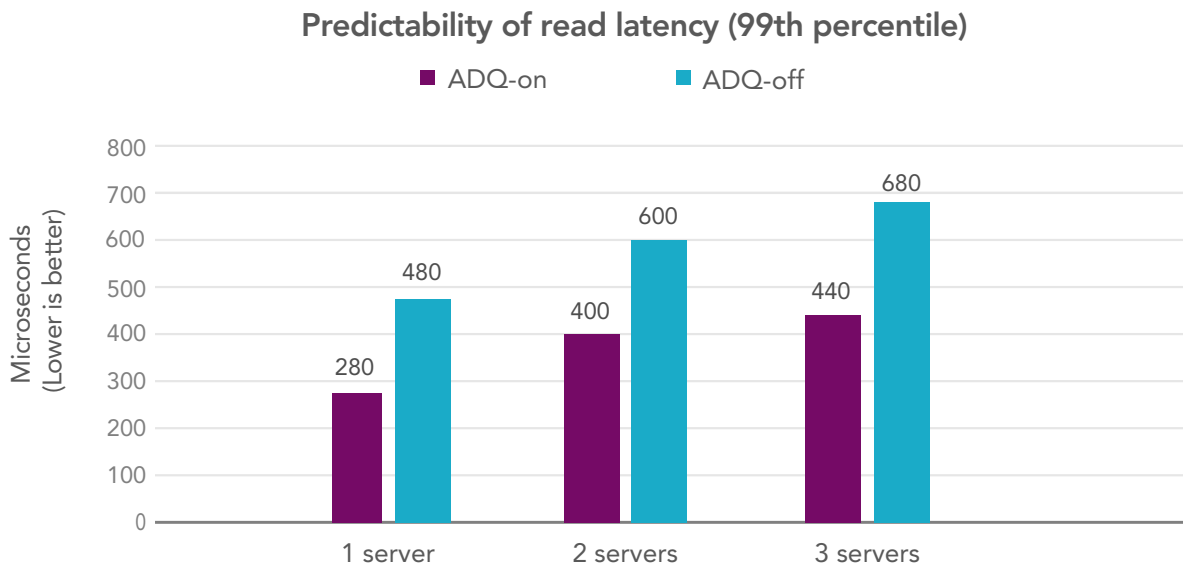

Figure 4: Latency for 99% of requests at multiple Dell EMC PowerEdge R740xd server node counts with Intel Ethernet 800 Series Network Adapters with ADQ on and ADQ off. Source: Principled Technologies.

Figure 5 on the following page presents a latency chart with histogram buckets comparing a single Dell EMC PowerEdge R740xd server node with ADQ enabled vs. the same server with ADQ off. The y-axis shows the percentage of requests, while the x-axis shows different latency thresholds in microseconds. For both solutions, 100% of requests were >40 microseconds. With ADQ enabled, there was a very low percentage of requests that took longer than 280 microseconds (as the line approaches 0), while with ADQ off over 20% of requests were longer than 280 microseconds. (Note: To see complete histogram charts for two- and three-node testing, visit the science behind the report.) By weighting the number of responses in each histogram bucket, we determined that enabling ADQ provided a 32.0 percent lower weighted average latency than a single server node with ADQ off.

Figures 4 and 5 show that with ADQ enabled, the Aerospike Database application experienced less variability and delivered most of its requests at a much lower latency. By ensuring fast response times as well as predictability, organizations can scale Aerospike databases with more confidence.

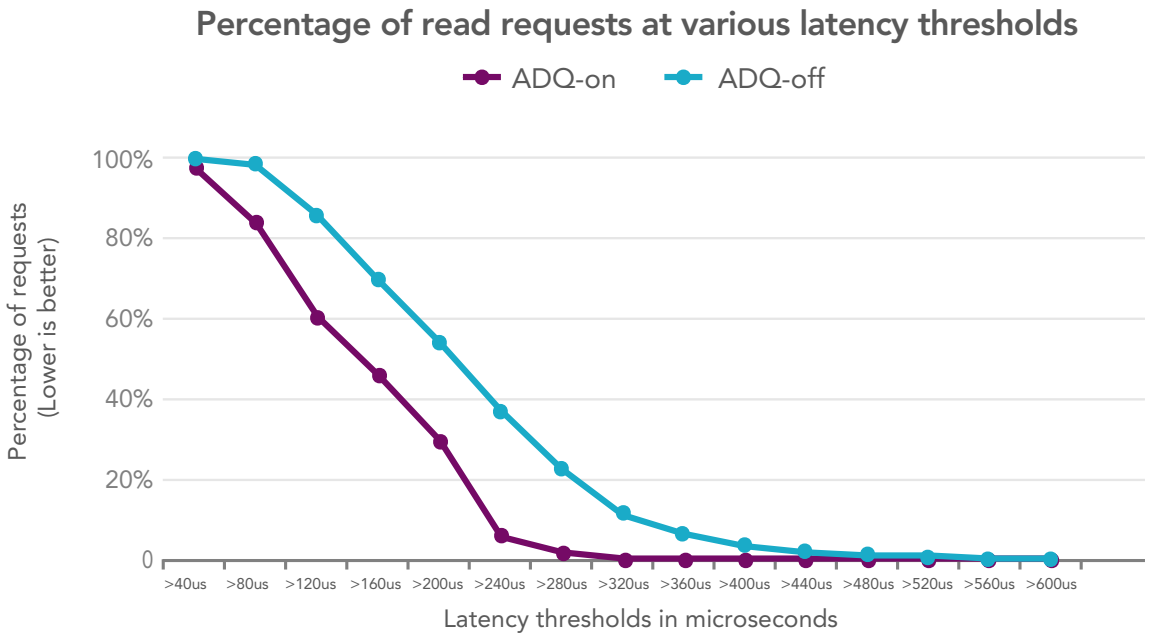

Figure 5: Latency histogram buckets for a single Dell EMC PowerEdge R740xd server node with Intel Ethernet 800 Series Network Adapters with ADQ on and ADQ off. Source: Principled Technologies.

About Intel Ethernet 800 Series Network Adapters

According to Intel, the new Intel Ethernet 800 Series offers technologies that "address a variety of workloads used in Cloud, Communications, and Enterprise market segments."[2] In addition to ADQ, 800 Series Network Adapters support two RDMA storage transport protocols, iWARP and RoCEv2, to offer storage flexibility for various business needs. To learn more about the Intel Ethernet 800 Series, visit https://www.intel.com/content/www/us/en/architecture-and-technology/ethernet/introducing-800-series.html.

To see how the iWARP and RoCEv2 RDMA storage transport protocols fared in our performance testing, visit http://facts.pt/gktwnar.

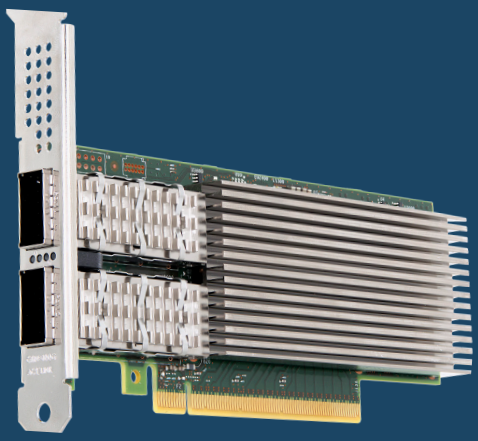

Enabling ADQ with Intel Optane persistent memory

After testing with traditional DRAM memory alone, we also compared the performance effects of enabling ADQ on a Dell EMC PowerEdge R740xd configured with Intel® Optane™ persistent memory (PMem). Servers that use 2nd Generation Intel Xeon Scalable processors, like the Dell EMC PowerEdge R740xd, offer support for Intel Optane PMem. As Figure 6 shows, enabling ADQ while using Intel Optane PMem showed performance gains similar to those we saw in our testing with traditional memory, increasing read throughput by up to 52.9 percent. Similar to DRAM testing, a single server node with Intel Optane PMem provided 29.1 percent lower weighted average latency with ADQ enabled than with ADQ off. This shows that whether you use Intel Optane PMem or DRAM, enabling ADQ can improve throughput, reduce latency, and increase response time predictability for Aerospike databases on the Dell EMC PowerEdge R740xd compared to leaving ADQ disabled.

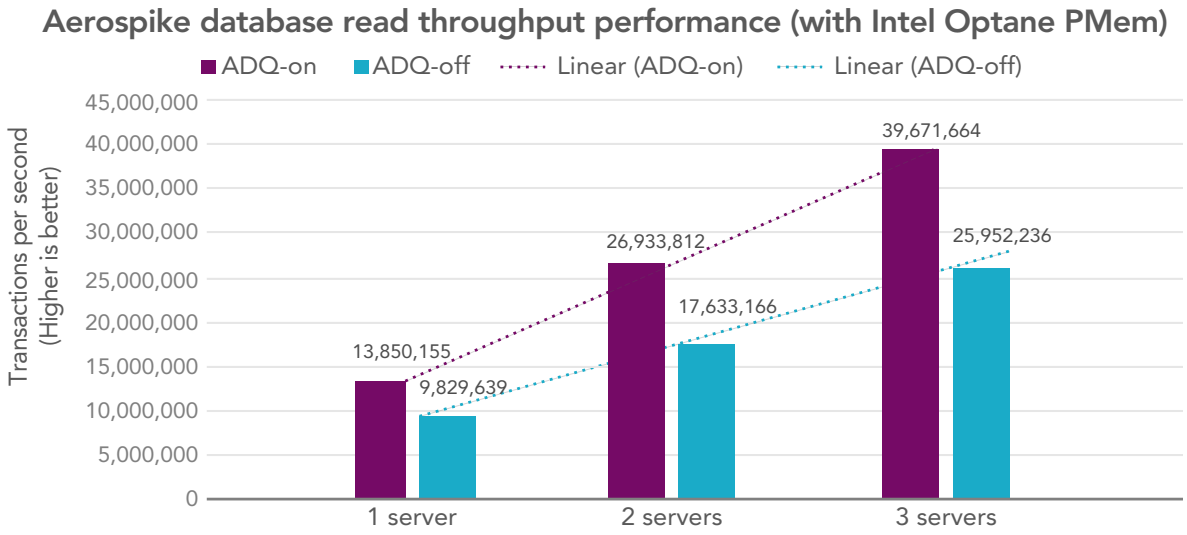

Figure 6: Aerospike database read throughput performance, in millions of transactions per second, of the Dell EMC PowerEdge R740xd at multiple server node counts with Intel Ethernet 800 Series Network Adapters with ADQ on and ADQ off using Intel Optane persistent memory. Source: Principled Technologies.

About Intel Optane persistent memory

Intel Optane persistent memory (PMem) is a new memory tier that offers both performance and persistence, making it an excellent choice for hot, frequently read data. Depending on your workload, you can use Intel Optane PMem in one of two modes:

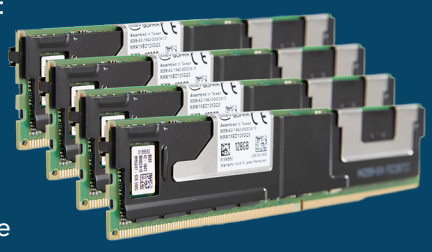

- Memory Mode, where Intel Optane PMem acts as large-capacity, volatile memory, with the system's DRAM acting as a cache. In this mode, latency is slightly higher, but capacity improves.
- App Direct Mode, where the applications and operating system are aware that two types of load/store memory are configured on the system, and applications can choose which type is best suited to meet the data's latency and persistence requirements.

In our testing, we used App Direct Mode. To learn more about Intel Optane PMem, visit https://www.intel.com/content/www/us/en/products/memory-storage/optane-dc-persistent-memory.html

Our tested server configurations with Intel Optane PMem provided comparable (a slight decrease) Aerospike performance to using DRAM alone, while offering additional benefits, including data persistence and a larger overall memory capacity per server node. Please note that in Aerospike testing, where you place the index and the database affect performance results; for Intel Optane PMem tests, we placed the index on DRAM and the database on PMem. For complete Intel Optane PMem testing results, including latency histograms, visit the science behind the report.

Conclusion

Leveraging Intel Ethernet 800 Series Network Adapters with Application Device Queues in Dell EMC PowerEdge R740xd servers can boost application throughput, lower latency, and increase predictability for your critical Aerospike Database applications. As our testing with traditional memory shows, enabling ADQ can help you make the most of your server resources by offering near-linear scaling and handling as many as 64.8 percent more transactions per second. Similarly, configurations with Intel Optane PMem improved read throughput by up to 52.9 percent when enabling ADQ. Turning on ADQ can help you shape your database traffic to deliver improved performance for users, up to 32.0 percent lower average latency, and 41.7 percent better P99 predictability, providing greater consistency meeting SLAs.

---

1 Aerospike, "Aerospike Database 5," accessed March 12, 2021, https://www.aerospike.com/products/database-platform/

2 Intel, "Move Data Faster," accessed September 24, 2020, https://www.intel.com/content/www/us/en/architecture-and-technology/ethernet.html

About the Dell EMC PowerEdge R740xd

What does the "xd" stand for? Extra drives—with up to 24 NVMe drives or 32 x 2.5" or 18 x 3.5" drives in just 2U of rack space, the Dell EMC PowerEdge R740xd promotes flexibility through several drive options and density that allows you to use software-defined storage solutions. The two-socket Dell EMC PowerEdge R740xd rack server features 2nd Generation Intel Xeon Scalable processors, embedded Dell EMC iDRAC9, and Dell EMC OpenManage™ software for management. To learn more about the Dell EMC PowerEdge R740xd, visit https://www.dell.com/en-us/work/shop/povw/poweredge-r740xd.

Read the science behind this report at http://facts.pt/8CTt1Ax ►

Principled Technologies®

Facts matter.®

This project was commissioned by Dell EMC.