

2X MongoDB Acceleration with 50% Higher Storage scalability

EXECUTIVE SUMMARY

MongoDB powers many business-critical applications across several industries, including retail, hospitality, travel, gaming, and financial services. It serves as an excellent data platform for organizations to build mobile applications, Internet of things, analytics applications, product catalogs, & content management. As the number of users accessing these applications increases and datasets grow from a few terabytes to hundreds of terabytes over a period, in order to sustain the read and write-intensive demands, MongoDB deployments have to scale from a few servers to 100s or even 1000s of servers with replicas and shards. It is essential to optimize the system and storage components as MongoDB scales to minimize the complexity and budget overspending. This brief focuses on the unique technical and financial benefits of deploying MongoDB with Pliops Extreme Data Processor (XDP) over traditional deployments.

Performance Scaling with Pliops

Pliops Extreme Data Processor (XDP) is a breakthrough data accelerator that delivers performance, capacity & reliability benefits while reducing the total cost of MongoDB deployment, operations, and scaling. It is a hardware-enabled key-value (KV)-based data storage engine in a PCIe card form factor designed to accelerate and scale data-intensive application workloads like MongoDB. It also includes built-in RAID5/6 style like data protection to tolerate any SSD failures improving MongoDB resiliency.

XDP can be implemented in any general-purpose server to accelerate and protect MongoDB using any NVMe SSDs. It is easy to deploy and provides a standard block interface for managing the data traffic from the MongoDB application to SSD storage devices. Pliops XDP accelerates storage software functions to optimize the management of data persistence and indexing tasks for NoSQL and databases while providing superior data compression and data protection benefits. Pliops XDP technology offers a simple and elegant way to manage data growth and it enables the user to economically scale MongoDB Clusters.

Pliops–MongoDB Solution Benefits

- Up to 2.3x improvement in performance acceleration for write-heavy workloads
- Elastic storage scalability to provide 50% higher storage capacity with built-in data compression engines
- Upto 2.4X reduction in latency to improve QoS of customer-facing applications
- Higher reliability and resiliency of MongoDB cluster to tolerate multiple SSD failures
- Reduce the infrastructure costs upto 50% without performance or quality of service (QoS) impact

MongoDB YCSB Performance Testing

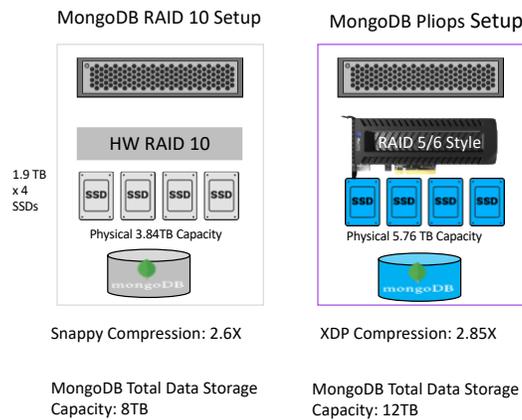


Figure1: 50% Capacity savings with XDP RAID & compression Offload

Yahoo Cloud Serving Benchmark (YCSB) was employed to evaluate the performance of MongoDB with and without Pliops setup. Both test setups included a 1.3TB dataset loaded using YCSB. As shown in figure 1, by switching from standard RAID 10 for MongoDB to Pliops Drive failure protection (DFP - RAID 5/6 style) and offloading compression to XDP by disabling snappy compression, we observed instant capacity savings of 50%. Total capacity per server increased from 9TB from the baseline scenario to 12TB with Pliops setup. Thus, Pliops XDP provides storage scalability that seamlessly scales with MongoDB data growth. The CPU savings with compression offload is discussed in the next section.

2X Performance Gains with Pliops

With the capacity expansion benefits realized using Pliops XDP setup using identical SSD storage devices, the next objective was to assess the performance benefits. The YCSB data loading, Write Heavy, Read Heavy, and Read-only benchmarks are utilized for performance testing evaluation. The results of running MongoDB with Pliops and traditional RAID 10 setup with Snappy for initial data loading are shown in Chart 1. Pliops setup provides 1.9X faster data loading by improving 25K data load operations/sec of SW baseline to close to 49.5K ops/sec and latency reduced from 0.31 millisecond (ms) to 0.15 ms which works out to about 51% latency reduction for 1.3TB dataset load.

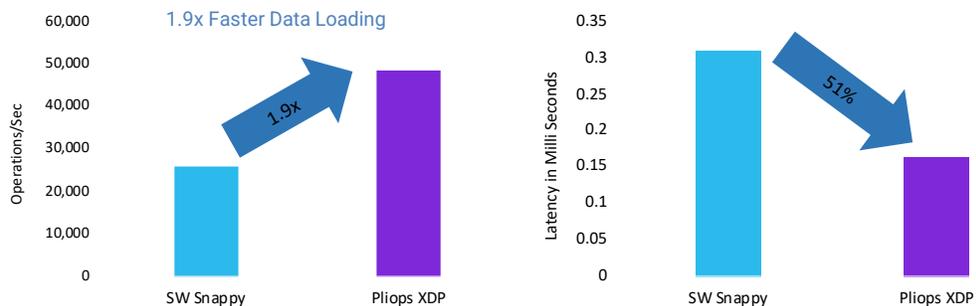


Chart 1: 1.9X Performance acceleration and 51% Latency reduction benefits for Data Loading with Pliops

For a mixed workload testing with 50% read and 50% write operations, as highlighted in figure 2 below Pliops XDP performance benefit over baseline increases to 2.3x, and performance increases from 31K ops/sec to 73K ops/sec. Thus, Pliops XDP increases MongoDB applications to quickly process incoming database write transactions and read queries simultaneously. In addition to the performance boost, the latency reduced from 4 ms to 1.6 ms, which is around a 2.4X latency reduction.

Latency reduction is another critical metric for application response time for improved user experience, especially during peak usage. Another significant benefit of Pliops XDP is eliminating the need for MongoDB default snappy compression and leveraging Pliops built-in compression that reduces the overall CPU utilization from 38% to 32%. Read-heavy and read-only workloads also benefited with 6 to 8% CPU savings, thereby freeing up CPU resources for increased user scalability and performance.

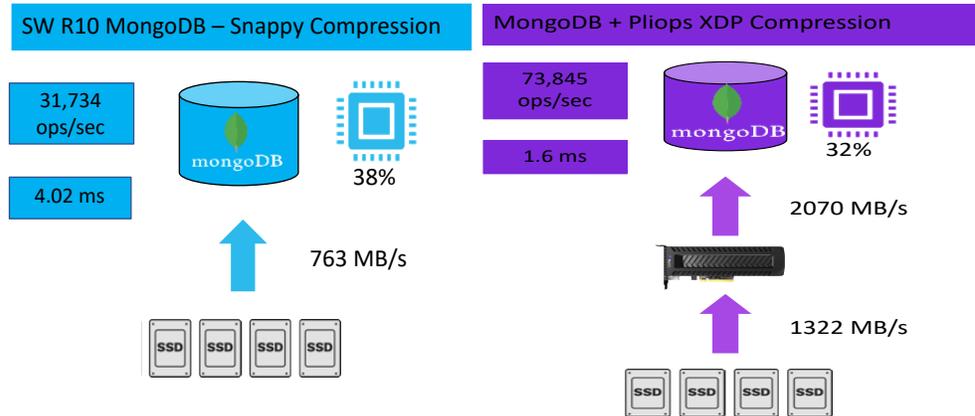


Figure 2: Mixed workload performance benefits of Pliops XDP

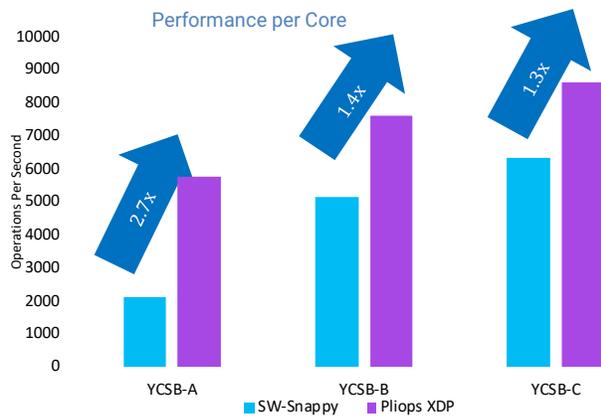


Chart 2: Mixed workload performance benefits of Pliops XDP

In addition to CPU savings, another important takeaway is performance per core benefit for SW baseline with Snappy compression and Pliops XDP with built-in data compression engines as shown in chart 2. The write heavy workload (YCSB-A) gets 2.7X higher performance per core. In the read-heavy (YCSB-B) workload, Pliops provides 1.4x benefit and 1.3x gain for Read-only workload.

Conclusion

Pliops XDP provides significant performance and latency benefits for economically scaling MongoDB applications from a few terabytes to many Petabytes. The built-in data compression, Pliops DFP feature, enables enterprises to efficiently manage the data growth challenges without performance and reliability impacts. This solution also provides significant cost savings by lowering the cost per terabyte and freeing up CPU resources for user scalability.

About Pliops

Pliops multiplies the effectiveness of organizations' infrastructure investments by exponentially increasing datacenter performance, reliability, capacity, and efficiency. Founded in 2017 and named as one of the 10 hottest semiconductor startups by CRN in 2020 and 2021. Pliops global investors include NVIDIA, Intel Capital, SoftBank, Western Digital, KDT, and Xilinx. **Learn more at www.pliops.com.**