



WHITE PAPER

Accelerating Big Data: Using SanDisk® SSDs for MongoDB Workloads

December 2014

SanDisk
a Western Digital brand

Western Digital Technologies, Inc.
951 SanDisk Drive, Milpitas, CA 95035

www.SanDisk.com

Table of Contents

Executive Summary3

SanDisk CloudSpeed® SSDs3

MongoDB3

YCSB benchmark4

Methodology5

 Datasets5

 Test Environment5

 Technical Component Specifications5

 Testing Configuration Details6

 MongoDB Configuration6

Test Workloads6

 Results Summary7

 Update Heavy7

 Read Heavy9

 Read Only10

Conclusion11

References11

Executive Summary

In recent years, as Big Data workloads have increased in the data center, NoSQL databases have become more widely used to store and access non-structured data. Examples of these data types include multimedia audio, video and photos, along with data from sensors from the Internet of Things. Solid state drives (SSDs) have shown their value as storage for these NoSQL databases by dramatically improving performance compared to mechanically driven hard-disk drives (HDDs).

To prove this advantage for SSDs, SanDisk tested MongoDB databases running on SSD-enabled server platforms. These tests show that HDDs often become the bottleneck in these NoSQL systems. Once the dataset size exceeds the memory capacity of the server, overall performance slows down. In contrast, SSDs improved the performance of the MongoDB NoSQL database workload. Importantly, we also looked at the impact on operational costs associated with data center space utilization, power and cooling costs—which decreased when SSD-based deployments were used.

SanDisk CloudSpeed® SSDs

SanDisk, a global leader in flash storage solutions, partners with all the leading storage vendors for meeting IT industry needs of flash-based products. The adoption of cloud computing is driving data growth, leading to an explosion in the volume of data that needs to be processed, stored and analyzed. These demanding Big Data workloads must be supported without compromising on performance, reliability, or longevity. SanDisk CloudSpeed SATA SSDs provide predictable performance and efficiency with superior reliability. These SSDs are secured by SanDisk's Guardian® capabilities of increasing durability, recoverability, and preventing data loss and corruption of data.



Figure 1: SanDisk CloudSpeed SATA SSDs

MongoDB

MongoDB is an open-source NoSQL document store database that is used in wide variety of workloads to support Mobility, Cloud Computing, Big Data/Analytics and other enterprise solutions. It provides application schema flexibility, which is not possible with relational databases. The relational databases (RDBMS) have schemas with tables and column attributes that have to be defined initially before loading the data for processing. MongoDB has a rich set of RDBMS functionalities like secondary indexes, query functionality and consistency for database transactions, but it does not require the up-front preparation for data-loading associated with RDBMS.

MongoDB provides scalability, performance and high-availability scaling from single server deployments to large, complex multi-site architectures. That gives it a broad range of deployment scenarios. It leverages in-memory computing (IMC) and it provides high performance for both reads and writes. MongoDB's native replication and automated failover features enable enterprise-grade reliability and operational flexibility.

Some of the important MongoDB features include:

- Data model: JSON data model with dynamic schemas
- Scalability: auto-sharding for horizontal scalability
- High availability: multiple copies are maintained with native replication
- Query model: Rich secondary indexes, including geospatial and TTL (Time-To-Live) indexes, aggregation framework and native map reduce
- Text search

YCSB benchmark

The Yahoo! Cloud Serving Benchmark (hereafter called YCSB) is a standard benchmark framework for evaluating the performance of new-generation cloud data serving systems like MongoDB, Cassandra, and Apache Hadoop HBase . The framework consists of a workload-generating client and a package of standard workloads .

YCSB evaluates the performance and scalability of cloud-based systems, while the performance section of the YCSB benchmark test focuses on measuring the throughput of the system for defined latency (delay in processing due to I/O data transfer) . Scalability focuses on the ability to scale elastically, so that these systems can handle more load as applications add more features, or ramp up to support an increased number of business users .

The YCSB benchmark also provides workload distribution options based on how real-time applications experience operations being requested of the system, such as insert/update/scan operations acting on a random set of data . YSCB workload distribution options are in two main “flavors,” as described here:

Uniform: This option for data-handling is based on assumption that all the records in the database are being uniformly accessed.

Zipfian: This is a statistical approach to handling requests to the database . Based on the assumption that the popularity of the record (e .g ., when World Soccer finals in Twitter is trending in popularity) is showing that it is being accessed more often than the other records in the database .

Latest: This option is based on the assumption that the “latest” events are more popular and are being accessed more frequently than the older events that have less frequent access.

Along with the workload distribution and the type of database operation being selected, the following workload types are used for this benchmark:

Workload	Operations	Record Selection/Distribution
Update Heavy	Read: 50% Update: 50%	Zipfian
Read Heavy	Read: 95% Update:5%	Zipfian
Read Only	Read: 100%	Zipfian
Read Latest	Read:95% Insert:5%	Latest

Figure 2: YCSB Workload

Methodology

The following sections describe the methodology that was used as the YCSB benchmark tests were conducted with both SSD-enabled and HDD-enabled servers supporting the MongoDB workload:

Datasets

The data was loaded to the MongoDB using the "load" phase of the YCSB benchmark tool .

Record Description: Each record consists of 10 character fields, each field 100 bytes long and Key assigned to each record which serves as a primary key.

Record Size: 1,024 Bytes

MongoDB Dataset Size: 32GB, 256GB, 1TB

Test Environment

The benchmark testing environment consists of one Dell PowerEdge R720 server with 24 Intel Xeon cores (two 12-core CPUs) with 96GB RAM used for hosting MongoDB server and one Dell PowerEdge R720 that serves as a client for YCSB benchmark tool. A 10GbE network interconnect is used between the MongoDB server and the YCSB client. The local storage is varied between hard disks (HDDs) and solid-state disks (SSDs) . The data set size of the YCSB tests was increased from 32GB, to 256GB and to 1 terabyte (1TB) . Figure 4 provides complete hardware and software components that were used for this testing environment.

Technical Component Specifications

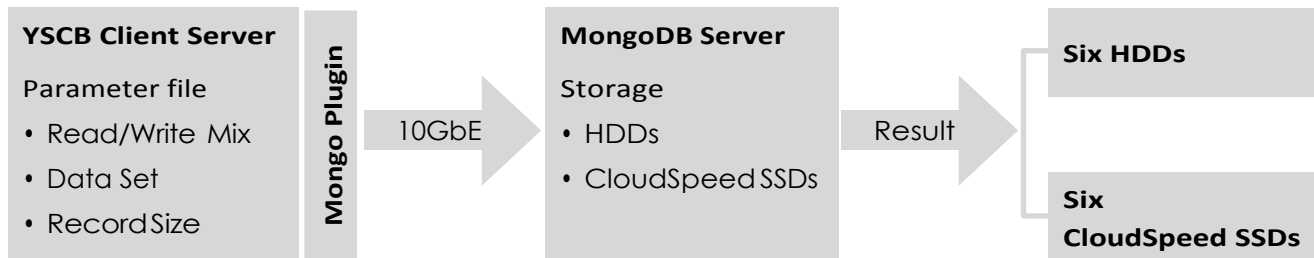


Figure 3: YCSB testing configuration

Testing Configuration Details

Hardware	Software if applicable	Purpose	Quantity
Dell Inc. PowerEdge R720 • Intel® Xeon® E5-2620 processor, two sockets, 24 cores (two 12-core processors) • 96GB memory	• CentOS 5.10, 64-bit • MongoDB : 1.2.2	MongoDB server	1
Dell Inc. PowerEdge R720 • Two Intel Xeon E5-2620 processor, two sockets, 24 cores (two 12-core processors) • 16GB memory	• CentOS 5.10, 64-bit • YCSB 0.1.4	YCSB client	1
Dell PowerConnect 2824 24-port switch	1GbE network switch	Data Network	1
500GB 7.2K RPM Dell SATA HDDs	Used as Just a bunch of disks (JBODs)	Data node drives	6
480GB CloudSpeed 1000 SATA SSDs	JBODs	Data node drives	6

Figure 4: Infrastructure details

MongoDB Configuration

MongoDB default configurations were used during the testing phase, and its data path and log path was switched between SSD and HDD for each testing cycle.

```
SSD Test: <MONGO_HOME>/bin/mongod -dbpath /sandisk/SSDDATA/mongodb/data -logpath /sandisk/SSDDATA/mongodb/log
```

```
HDD Test: <MONGO_HOME>/bin/mongod -dbpath /sandisk/HDDDATA/mongodb/data -logpath /sandisk/HDDDATA/mongodb/log
```

Test Workloads

The primary objective of this benchmark test was to identify the advantage of using SanDisk SSDs for a MongoDB NoSQL store, and to provide performance data points for SSDs and HDDs. This benchmark consists of single-node MongoDB database with the standard YCSB benchmark workload types A, B and C, and a plan to test it with three different dataset sizes.

- The YCSB workload types based on percentage of reads and writes:
 - Workload A: Update Heavy: 50% Update / 50% Read
 - Workload B: Read Heavy: 5% Update / 95% Read
 - Workload C: Read Only: 100% Read Only
- The YCSB default data size: 1 KB records (10 fields, 100 bytes each, plus key)
- Size of the data set is 200,000 key/value pairs
- The data set types are as follows:
 - In-memory dataset: 32G
 - Disk dataset 1: 256G
 - Disk dataset 2: 1TB
- The YCSB workload distribution types are as follows:

- Uniform: All database records are uniformly accessed
- Zipfian: Some records in the database are accessed more often than other records

Results Summary

Based on test results, from an operations-per-second perspective, the MongoDB performance on solid state disks (SSDs) is outstanding compared to hard disk drives (HDDs) for the same MongoDB configuration. This advantage gets further highlighted when the dataset goes beyond the memory capacity of the MongoDB server. The latency metrics for SSDs, for both the read and write operations, were the lowest across all workloads –which is an important factor regarding the scalability of MongoDB server.

Update Heavy

In-memory dataset: Figure 5 shows update-heavy workload results for the 32GB dataset, which is smaller than the memory capacity of the MongoDB server. SSD performance has higher throughput for both the Uniform and Zipfian workload types.

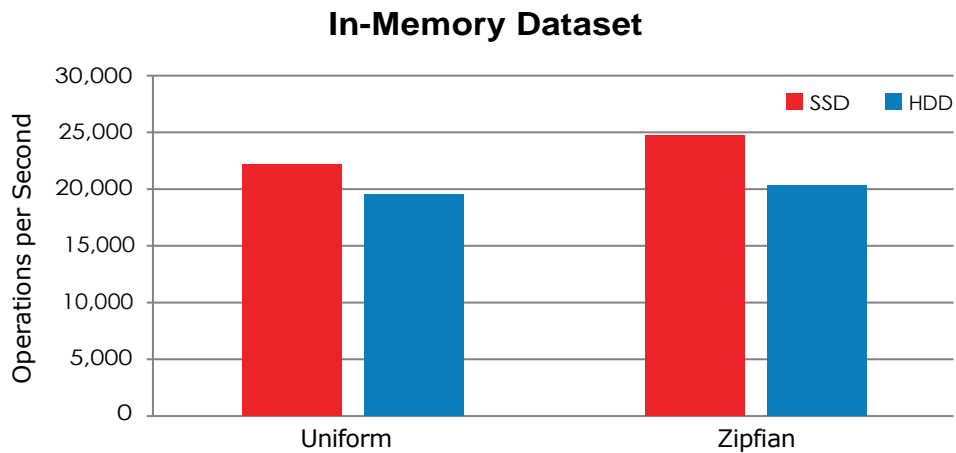


Figure 5: Throughput comparisons of update-heavy in-memory dataset

On-disk dataset: Dataset results for the 256GB and 1TB database sizes. These two data sets exceed the capacity of available memory and must reside on an HDD. As seen in Figure 6, SanDisk SSD performance is far superior compared to HDD performance in this scenario.




Figure 6: Throughput comparisons of update-heavy on-disk dataset

YCSB Workload Types	Storage Configuration	YCSB Workload Distribution	32GB	256GB	1TB
Workload A (50r/50w)	HDD	Uniform	19,490	95	66
Workload A (50r/50w)	HDD	Zipfian	20,300	165	107
Workload A (50r/50w)	SSD	Uniform	22,124	2,418	1,871
Workload A (50r/50w)	SSD	Zipfian	24,732	4,676	3,523

Figure 7: Throughput results of update-heavy workload

Latency: SanDisk CloudSpeed SSDs provide consistently low latency results, even with large datasets for both read and write operations .




Figure 8: Latency results for SSD vs. HDD update-heavy workload

Read Heavy

SSDs provide excellent performance results for read-heavy workload . As the data set expands from 32GB to 256GB to 1TB, the SSD advantage gets clearly highlighted as shown in Figure 9 (on-disk dataset) .




Figure 9: Throughput results of read-heavy workload

Latency: SSDs deliver minimal latency for read-heavy workload and this advantage is more pronounced for large datasets (256GB and 1TB) as shown in Figure 10 and for some datasets HDDs generate up to 2.3x higher latency.




Figure 10: Latency results for read-heavy workload

Read Only

This workload is exclusively a read-only workload, which is fetching large amounts of data from MongoDB server. As expected, the SSD gets clear advantage in this workload as the size of the data set exceeds available memory in going from 256GB to 1TB.




Figure 11A: Throughput results for read-only workload




Figure 11B: Latency results for read-only workload

Latency: SSDs deliver virtually no latency for in-memory datasets, and for large datasets beyond memory size, SSDs delivers minimal latency. HDDs for large datasets encounter up to 43x higher latency, highlighting the benefit of using SSDs for such large workloads.

Conclusion

SanDisk CloudSpeed SSDs deliver superior performance throughput—and they do so with a consistently low latency—for all the workload and dataset types . This kind of platform, with high performance and low latency platform, helps the MongoDB database to complete all of its operations in shorter time intervals than it would with HDDs, thereby reducing the need for number MongoDB servers in a given clustered-server environment . Reduction of MongoDB cluster density reduces both capital expenses (CAPEX) and operational expenses (OPEX), with fewer MongoDB database instances to manage and administer. SanDisk’s Guardian technology, which ships with the SanDisk SSDs, provides a data protection capability, thereby securing the customer’s investment in these solid state disks.

References

SanDisk website, includes information about Guardian technology: www.sandisk.com MongoDB:

www.MongoDB.org

Yahoo LABS Yahoo! Cloud Serving Benchmark:

http://research.yahoo.com/Web_Information_Management/YCSB/

MongoDB memory mapping mechanism:

www.polyspot.com/en/blog/2012/understanding-mongodb-storage/

Specifications are subject to change. ©2014 - 2016 Western Digital Corporation or its affiliates. All rights reserved. SanDisk and the SanDisk logo are trademarks of Western Digital Corporation or its affiliates, registered in the U.S. and other countries. CloudSpeed, CloudSpeed Eco, CloudSpeed Ascend, CloudSpeed Ultra, Guardian Technology, FlashGuard, DataGuard and EverGuard are trademarks of Western Digital Corporation or its affiliates. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s). 20160623

Western Digital Technologies, Inc. is the seller of record and licensee in the Americas of SanDisk® products.