WHITEPAPER

Improve Performance of MongoDB with SanDisk ZetaScale™ Software

SanDisk®
951 SanDisk Drive, Milpitas, CA 95035
©2014 SanDisk Corporation. All rights reserved
www.SanDisk.com

Executive Summary

In its whitepaper, *Accelerating Enterprise Applications and Reducing TCO with SanDisk ZetaScale™ Software*, SanDisk demonstrated that ZetaScale can improve performance and costs for a variety of applications.

This whitepaper demonstrates how enhancing MongoDB with SanDisk ZetaScale technology will improve the performance of MongoDB by up to 3x over SSDs alone and 20x over HDD.

Introduction

MongoDB is an open source NoSQL database. Unlike traditional relational databases, MongoDB enables unstructured or schema-less databases. As such NoSQL databases are highly scalable – and can handle large data files. Importantly, they can address multiple data-types, including structured, semi-structured and unstructured data—with all data stored in a single data repository.

The performance of MongoDB will improve significantly when flash storage is used to replace traditional HDD. Using ZetaScale with flash will further improve this performance and also lead to server consolidation and lower operating costs due to lower power consumption of SSD storage.

MongoDB Overview

MongoDB is a NoSQL document store designed around the following features:

MongoDB stores data in JSON documents (serialized to BSON), providing a rich data model that seamlessly maps to native programming language types. The dynamic schema makes it easier to evolve a program's data model compared to enforced schemas such as an RDBMS.

MongoDB provides many of the features of a traditional RDBMS such as secondary indexes, dynamic queries, sorting, rich updates, and easy aggregation. This provides the breadth of functionality of an RDBMS, with the flexibility and scaling capability that the non-relational model allows.

By keeping related data together in documents, queries can be much faster than in a relational database where related data is separated into multiple tables and then needs to be joined later. MongoDB also makes it easy to scale databases by autosharding, which allows you to scale your cluster linearly by adding more machines. It is possible to increase capacity without any downtime, which is very important for web servers when load can increase suddenly and bringing down the website for extended maintenance can negatively impact revenue.

ZetaScale Technical Overview

SanDisk ZetaScale software provides an object store subsystem for use by applications (or databases).

Applications can place and retrieve their logical objects into the ZetaScale key-value store. ZetaScale optimizes the use of CPU cores, DRAM, and flash to maximize application throughput. The library of flash-optimized software code can provide performance levels that are close to in-memory performance speeds.

ZetaScale significantly boosts database and application performance, not only compared to HDD-based systems, but also compared to standard flash-based systems.

ZetaScale utilizes the unique features of flash, providing about 3x more throughput than non-optimized applications running on identical flash-based hardware configurations.

ZetaScale's highly parallelized code supports high throughput for flash I/Os, due to its support of fast operations, with high rates of I/Os per second (IOPs), even for small objects.

ZetaScale automatically and efficiently maps object keys to storage locations. ZetaScale automatically manages object migrations within the DRAM-flash storage hierarchy to maximize performance, including using DRAM as a cache to the flash.

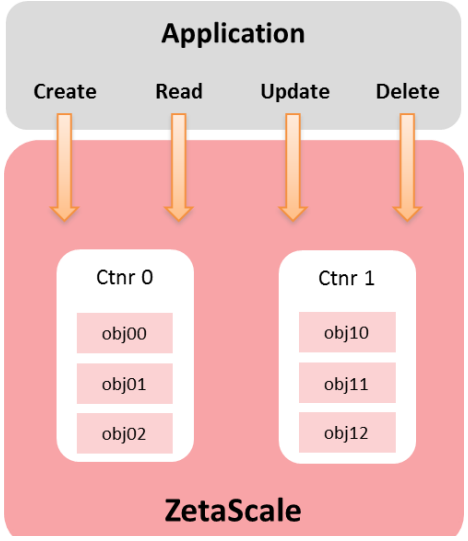

Figure 1: SanDisk ZetaScale Software architecture overview

MongoDB enhanced with ZetaScale Software

The power and performance of MongoDB are best when the dataset is in memory. Operations where documents need to be created and retrieved from storage will provide a significant drop in performance. This can be improved by allowing MongoDB to directly call the ZetaScale API's for optimized throughput.

Also ZetaScale provides built in logging and persistence so the MongoDB journaling can be avoided.

About the test

The benchmark test is set up to simulate a realistic working environment, as follows:

Server

- 24 cores Intel Westmere with 96GB DRAM

Storage

- 8x OCZ with SW RAID0

Workload

- YCSB benchmark
- 128GB 1KB objects

Software

- ZetaScale v1.2
- MongoDB v2.2.0

Test Results

This setup was tested by applying the following workloads:

1. Workload A: 100% Read
2. Workload B: 95%/5% Read/Write

The performance for different workloads is illustrated below:

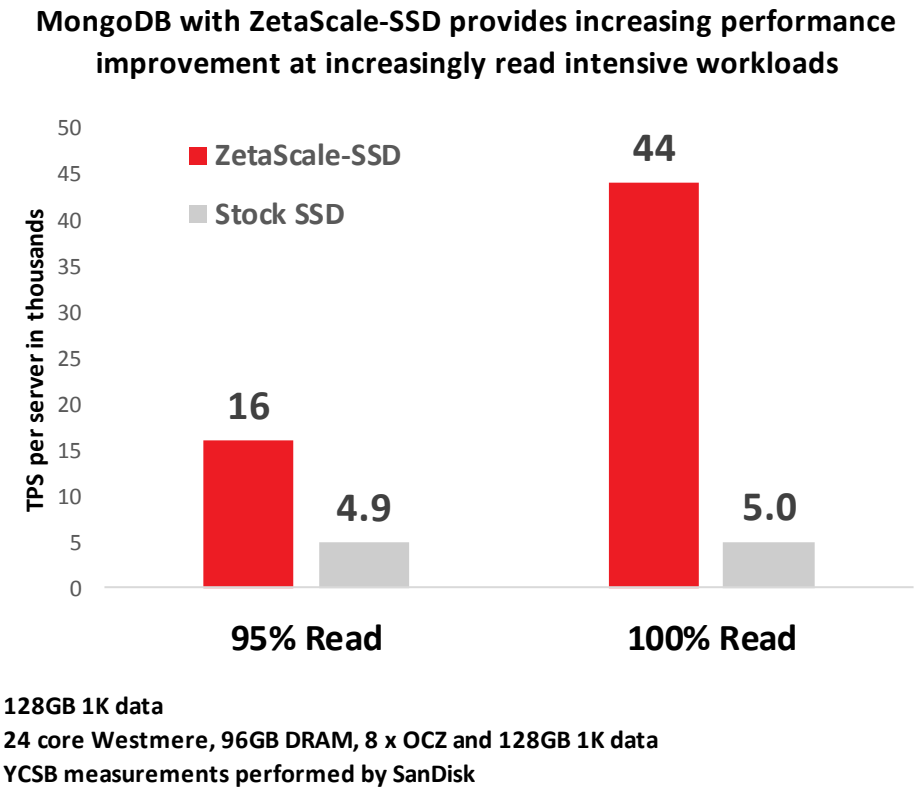

Figure 2: MongoDB performance improvement by using ZetaScale software

Figure 2 clearly demonstrates that the performance of MongoDB with ZetaScale and flash is more than 3x better for a 95%Read / 5%Write workload and up to 9x better for a Read-only workload.

This results in a 20x performance improvement over HDD, as illustrated below:

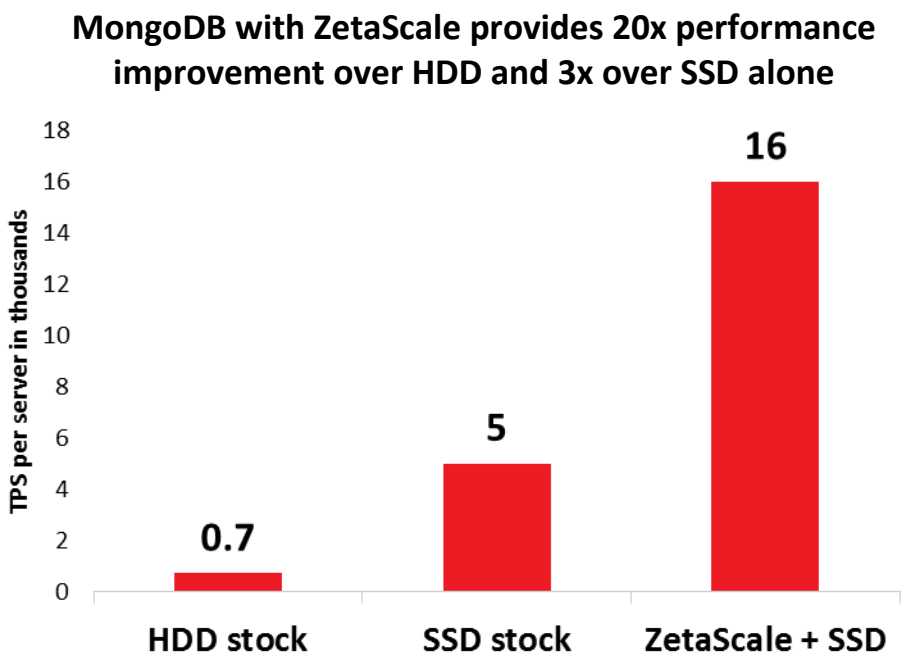

Figure 3: TPS per server for stock MongoDB with HDD and standard SSD compared to ZetaScale with SSD.

Conclusions

This test clearly demonstrates that MongoDB with ZetaScale and flash performs at a significantly improved level compared to stock MongoDB at different read/write levels.

ZetaScale utilizes the unique features of flash and optimizes the use of CPU cores, DRAM and flash to maximize application throughput.

©2014 SanDisk Corporation. All rights reserved. SanDisk is a trademark of SanDisk Corporation, registered in the United States and other countries. ZetaScale is a trademark of SanDisk Enterprise IP LLC. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their holder(s).