LSI
Storage. Networking. Accelerated.

White Paper

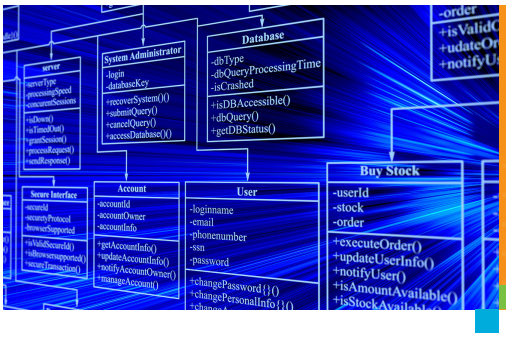

Accelerate Database Performance and Reduce Response Times in MongoDB "Humongous" Environments with the LSI® Nytro™ MegaRAID® Flash Accelerator Card

Summary

One of today's growing database types is MongoDB in which applications such as real-time monitoring or content management may be performed.

The flash caching algorithm in Nytro MegaRAID technology allows hot data in large collections of databases to be quickly recognized as the data set changes helping to provide consistent improvements in performance.

In tests conducted by LSI, the implementation of MongoDB installations using a Nytro MegaRAID flash acceleration card doubled operations per second and reduced average latency by up to 60% for a uniform distribution of data.

The Rise of MongoDB

MongoDB, derived from "humongous", is an open-source document oriented database and the leading NoSQL database system. A powerful, flexible and scalable data store that provides high performance, high availability and easy scalability, MongoDB has features that rival relational databases including: a shard-capable scale-out model, primary, secondary and geospatial indexing and built-in support for MapReduce. MongoDB is used by thousands of companies worldwide to provide services for applications such as content management, Platform as a Service (PaaS), real-time monitoring, CRM, custom analytics stores, and many more. The MongoDB data model allows it to automatically split up data across multiple servers providing easier scalability and high availability. By adding additional nodes to the cluster, administrators can easily scale capacity and performance. High availability (HA) is also achieved through its automatic failover feature, in which a backup slave can automatically be promoted to a master in the event of a failure.

With MongoDB, a document is the basic unit of data (equivalent to a row in a RDBMS), a collection consists of one or more documents (similar to a table in an RDBMS) and a database consists of one or more collections. A single instance of MongoDB can host multiple independent databases. Documents of the same kind are typically grouped together in the same collection to provide data locality and create excellent opportunities for data tiering and caching. As the data set size grows, it can be automatically sharded (the process of storing data records across multiple machines; MongoDB's approach to meeting data growth demands) across the cluster for scalability. In addition, documents can be batch inserted (and removed) with minimal overhead, therefore, it is important to build an optimal storage subsystem to assure the fast loading of data. As is typical in most databases, there are regions of frequently accessed data that would benefit from low latency flash caching. MongoDB by itself does not include caching or tiering features.

Accelerate Performance in MongoDB Environments with the Nytro MegaRAID Flash Accelerator

Testing conducted by LSI compared a six-disk RAID 10 configuration with and without flash-based cache located on the disk controllers of each node in a 3-node replica set. An additional node was used as the workload driver so that the overhead associated with workload generator activity did not interact and introduce measurement noise. Each node ran CentOS 6.4 configured with 16GB of RAM, and the Nytro MegaRAID cache size was 44GB configured as a RAID 1 to ensure the highest data availability. Using a version of the Yahoo Cloud Serving Benchmark (YCSB) forked specifically for MongoDB (version 0.14) to measure database performance, the client was configured with the default write concern which acknowledges writes to the primary server.

The LSI Nytro MegaRAID card offers MegaRAID data protection and flexible onboard flash caching technology that can be used in a variety of ways to help improve performance and server storage density.

Existing server deployments currently using a standard RAID controller or HBA have the potential for higher performance simply by replacing the existing card with a Nytro MegaRAID card.

All reads and writes from YCSB occur on the primary node. The database contains a single collection with 75 million documents and is 100GB in size. Additional workloads (workloads B and F) were also used. Workload B is primarily read-based (95% reads and 5% writes) and simulates applications such as photo tagging or status updating. Workload F reads a record, modifies it and then writes it back, simulating an I/O pattern found in transactional type databases.

In order to measure steady-state performance, testing was performed in six consecutive 10-minute phases to obtain the IOPS and latency stats for average, 95th percentile and 99th percentile over each interval. This approach measures the impact of "warming up" the Nytro MegaRAID flash cache.

Test Results

The latency and operations per second statistics are shown in the charts below. Using the Nytro MegaRAID card's automatic flash caching technology, the number of achieved operations per second improved by 117%, while the response times were reduced by approximately 50%. The warm up time (the amount of time it took to identify and begin caching hot data), was less than 10 minutes. This improvement applied to both workload profiles.

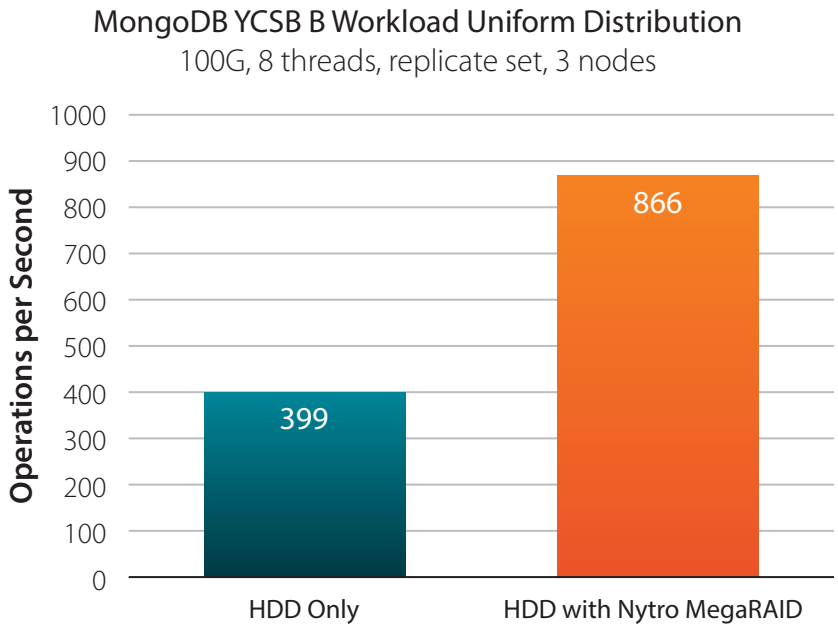

The following chart compares the average read latency for uniform distribution for flash cache-accelerated and hard disk drive (HDD) cases:

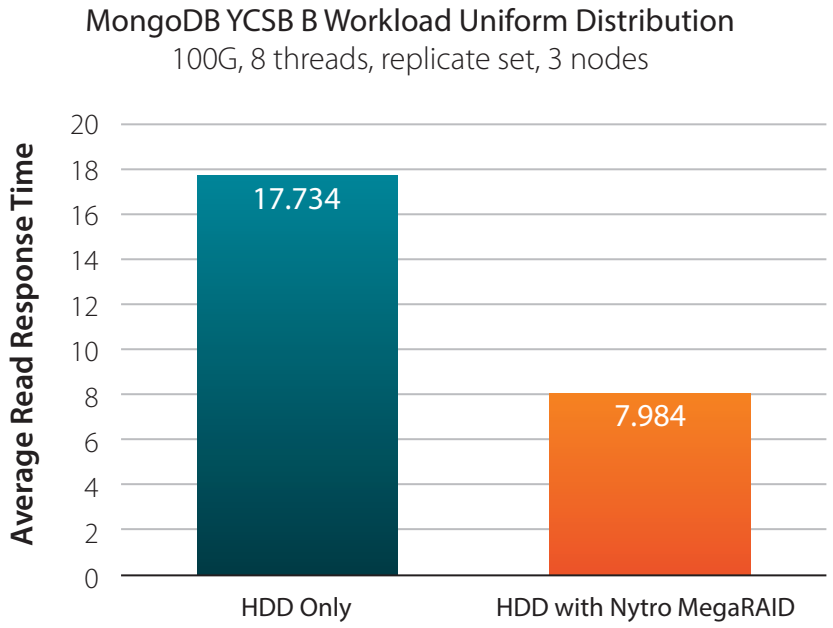

The following chart compares the average update latency for uniform distribution for flash cache accelerated and HDD cases:

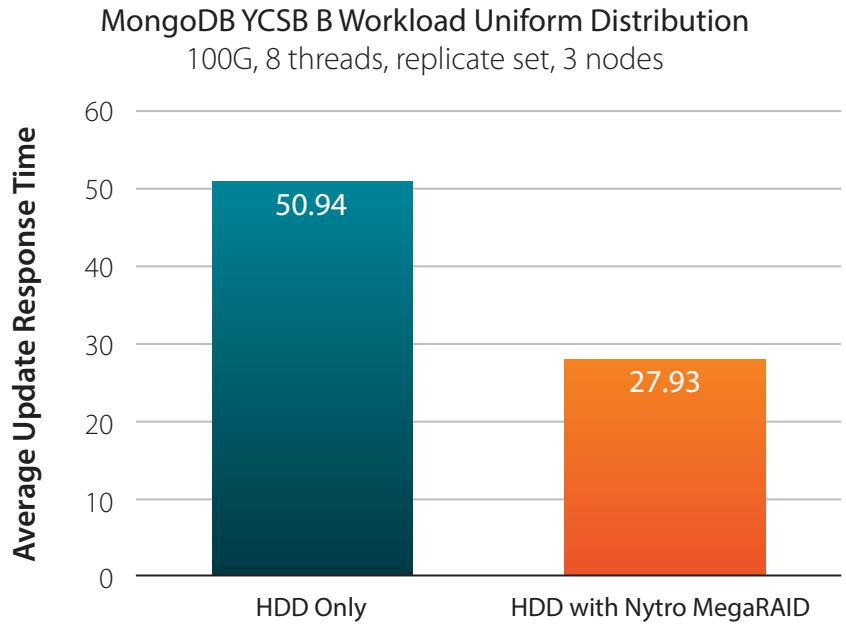

The following chart compares the operations per second for uniform distribution for flash cache-accelerated and HDD cases:

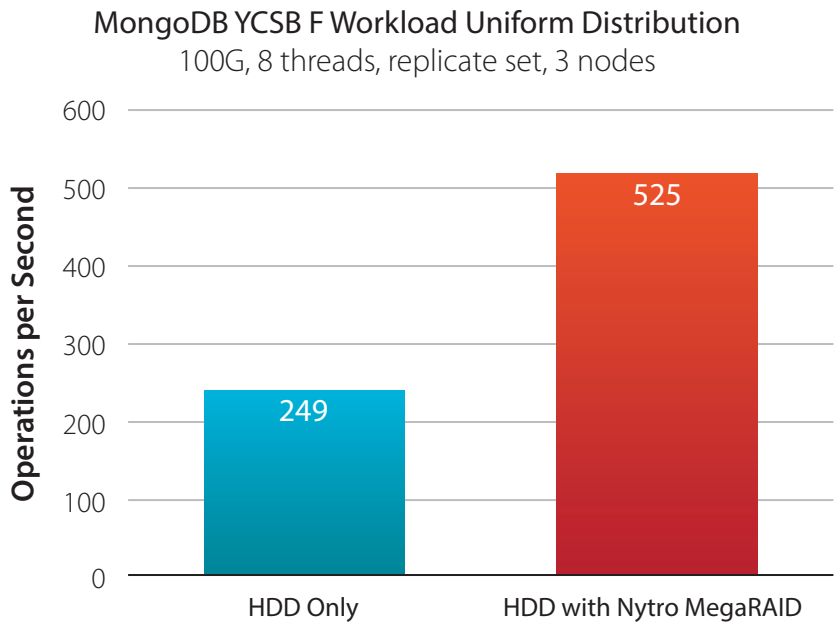

The following chart compares the average read latency for uniform distribution for flash cache-accelerated and HDD cases:

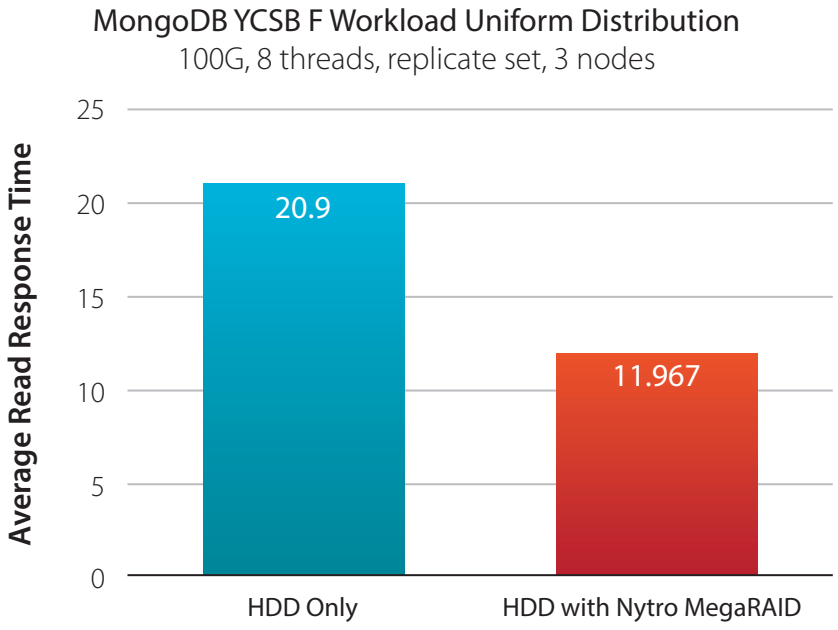

The following chart compares the average update latency for uniform distribution for flash cache-accelerated and HDD cases:

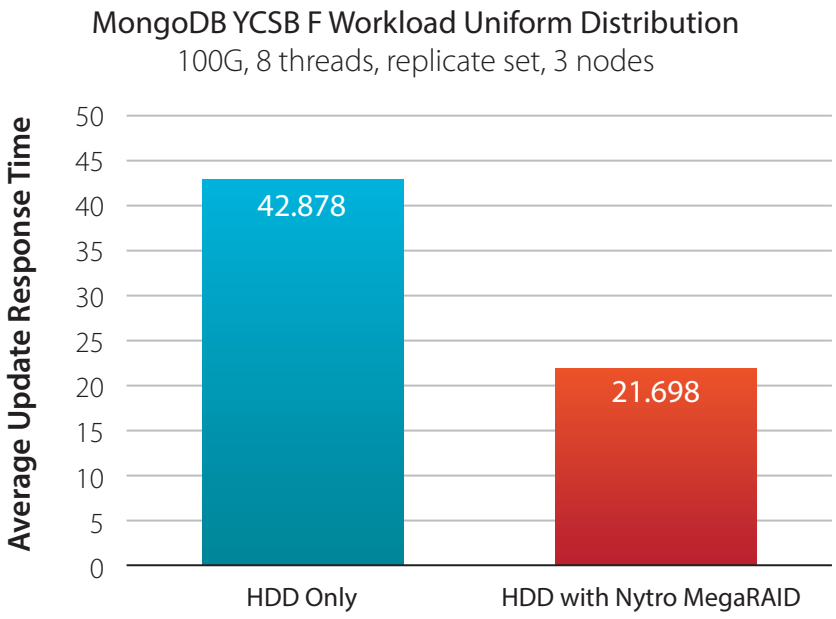

Conclusion

Today's megatrends – big data and cloud computing - are driving the adoption of NoSQL technologies as a viable alternative to relational databases when operating at scale. The ability to store schema-less data is an attractive solution for the variety, velocity and volumes of data captured and processed today. These YCSB test results clearly show that flash caching is a valuable and cost-effective way to accelerate MongoDB installations. The table below summarizes the test results as factors of improvement in performance for operations per second and latency.

	Workload B (Web Updates)	**Workload F** (Transactions)
Operations Per Second (% Improvement)	117%	111%
Average Read Response Time (% Reduction)	45%	50%
Average Update Response Time (% Reduction)	55%	43%

The LSI Nytro MegaRAID card's effective caching algorithm allows the hot data to be recognized quickly as the data set changes helping to provide consistently improved performance. Data management structures such as indices can be retained in flash cache to allow even faster lookups and efficient access of document records.

Test Disclosure

- MongoDB v2.4.8
- JDK 1.6.45 (Oracle)
- OS: CentOS release 6.5
- 4 machines, 1 workload driver, 3 replica nodes
  - 16 core, 64bit, 16GB RAM
  - Intel® Xeon® CPU E5-2650 0 @ 2.00GHz stepping 07, SandyBridge
  - 10gb dedicated Ethernet
- LSI Nytro MegaRAID flash accelerator card
  - 88GB flash based write back cache
  - 6x 1TB 7200 RPM SATA disks
- YCSB (Yahoo Cloud System Benchmark) 0.1.4
  - 75,000,000 records
  - 8 threads

References

https://github.com/brianfrankcooper/YCSB/tree/master/mongodb

http://www.mongodb.org/

LSI®
Storage. Networking. Accelerated.®

For more information and sales office locations, please visit the LSI website at: www.lsi.com

North American Headquarters
San Jose, CA
T: +1.866.574.5741 (within U.S.)
T: +1.408.954.3108 (outside U.S.)

LSI Europe Ltd.
European Headquarters
United Kingdom
T: [+44] 1344.413200

LSI KK Headquarters
Tokyo, Japan
T: [+81] 3.5463.7165