

Business DBs Demand All Flash Storage

An SSD (Solid State Drive) can deliver over 1000 times more performance than an HDD (Hard Disk Drive) depending on the application, NAND type, system, SSD interface, SSD model and manufacturer. Databases are performance hungry environments and are usually mission and business critical. A poorly performing database can negatively impact employee efficiency, customer satisfaction and revenue generation. It's no wonder more and more companies are using all flash storage to host their database applications. Databases need to perform or the entire business suffers.

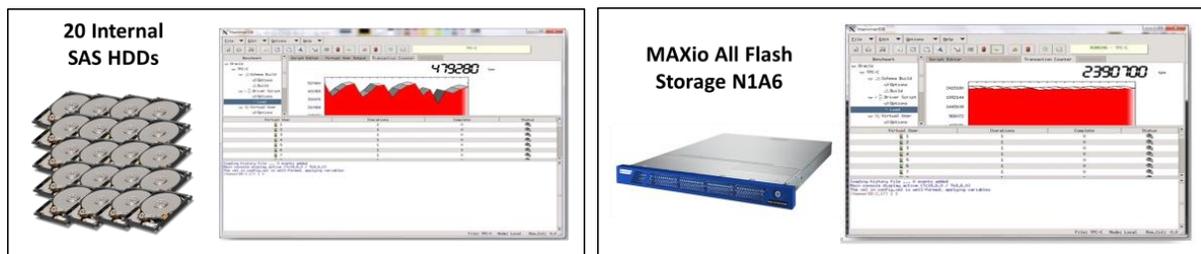
Hosting a database on a hard drive based system is usually out of the question. On average, HDDs deliver up to 400 IOPs while SSDs deliver up to an astounding 450,000 IOPs of random reads. This has an enormous impact to database performance.

We recently put together an Oracle performance benchmark to compare a server with 20 internal SAS RAID'ed HDDs to a 20-drive external iSCSI connected SATA SSD RAID storage system. We used HammerDB. HammerDB is a graphical open source database load testing and benchmarking tool for Linux and Windows to test databases running on any operating system. HammerDB is automated, multi-threaded and extensible with dynamic scripting support.

HammerDB includes complete built-in workloads based on industry standard TPC-C and TPC-H benchmarks as well as capture and replay for the Oracle database. HammerDB includes transaction and CPU monitors to complement the rich feature set that makes HammerDB the leading tool for benchmarking, testing and comparing the leading databases worldwide.

The 20-drive HDD system managed fewer than 480,000 transactions per minute, approximately 400 transactions per second per drive. While the 12TB SSD system, using mid-range 2.5 SATA SSDs that were saturated, managed nearly 2,400,000 transactions per minute. The SSD system could easily manage an Oracle DB workload which is 5 times larger than what the HDD system could support.

(NOTE: It is well known that SSDs lose performance overtime. An SSD saturated with data will be much slower than a new SSD. Most midrange SATA SSDs use trim and garbage collection technologies to maintain maximum performance. Trim allows an operating system to inform a solid-state drive (SSD) which blocks of data are no longer considered in use and can be wiped internally. Garbage collection works in the background with trim and systematically clears these blocks of data during off-peak times.)



A quick note about hybrid systems for database applications

Hybrid systems use SSDs as cache and HDDs to store the majority of the data. This makes write operations more performant. Read operations are also more performant since recently written or frequently accessed data is kept on the SSD cache. But only the most expensive hybrid systems support this feature. In many other hybrid systems the SSD cache needs to be defined as a separate volume. Write and read caching now requires special virtualization software and a great deal of management by the IT staff. There is also the problem of a cache miss



where the database application needs a certain block of data that is not on the SSD cache. This greatly degrades performance since the data needs to be retrieved from the much slower HDDs.

In order to have high and predictable performance, more and more companies are adopting an all flash array solution over a hybrid solution. And with the prices of all flash storage dropping (now only \$5 to \$6 per usable GB) it's much easier to justify all flash storage for mission and business critical database applications. Moreover these systems are easy to configure and manage.

Supporting Business and Mission Critical Databases requires 24 x 7 x 365 System Availability

We live in a 24 hour a day world where data must be available around the clock. Mission and business critical databases need to be online to service a global community that needs access to data to complete their tasks. Other applications, like data analytics, need the same level of uptime for segments such as security, trading, transportation, travel, emergency medical, national defense, social media, entertainment...the list goes on and on. Databases need to be continually available to service the needs of clients and employees worldwide. Therefore the infrastructure and storage systems they rely on need the same level of 24 x 7 x 365 availability.

Some All-Flash Arrays, like those from BITMICRO, are now featuring complete high availability. There is no single point of failure, everything is redundant: network ports and adapters, SSD drives, power supplies, system controllers. In fact even the storage is redundant. There are two independent data sets on separate SSD RAID sets, each has a separate controller. The data is mirrored using InfiniBand. Even if an entire enclosure would go off line, the data would still be available. The only thing that is shared is the physical rack mountable chassis and backplane.

Moreover this level of high availability provides for more predictable performance. As an example, the database application does not have the added burden of providing protection such as mirroring data to ensure business continuance in the event a storage enclosure suffers a catastrophic failure. The database application can be focused on performance. Second, because any failure does not impact performance, user response times are always consistent even during a component failure.

In addition all flash arrays provided by BITMICRO offer software RAID and performance enhancing technology. All data and parity are evenly distributed across all the SSDs in the RAID set. To maintain performance during a drive rebuild, the SSDs are divided into two groups. If an SSD fails in the first group, this group will not accept any writes. All the writes and associated parity will be directed to the second group of SSDs within the enclosure. Although both SSD groups will continue to respond to read requests, only the SSD group not restoring a drive will be given writes. This provides the group restoring the drive with more resources to quickly bring the replacement SSD online with virtually no impact on performance.

HDD arrays are being replaced by SSD All-Flash Arrays to meet the ever increasing demands of performance hungry databases. There is no better time than now to discover the operational benefits of using solid state storage to accelerate your business, improve employee efficiency and greatly increase customer satisfaction.

Writer: Zophar Sante, Business Development

Date: 12/12/2016

For more information on solid state storage systems visit: www.bitmicro.com MAXio® All Flash Storage