# LSI® Nytro™ XD Caching Solution Evaluation

***Evaluation report prepared under contract with LSI Corporation***

## Introduction

IT professionals who oversee the performance of mission-critical applications are looking for ways to increase performance and decrease latency for these applications. While these IT professionals want to improve the performance of these applications, they do not have unlimited resources and must achieve these improvements frequently with limited budgets.

Enterprise-level flash technology is transforming the computing environment and is bringing high performance storage to the datacenter environment. Flash-based caching is an excellent way to bring the high performance and low latency of flash technology at a minimal cost and without requiring application changes or back-end storage changes.

Read-intensive workloads can be very time-consuming and stressful on a storage system. For these workloads, either there is a fixed amount of work that must be completed as fast as possible, or there is an ongoing amount of bursty work with relatively unpredictable load levels. In either case, performance and latency are critical for these workloads, and anything that can be done to increase performance and decrease latency improve the user experience and allow work to be completed in less time.

LSI commissioned Demartek to evaluate its Nytro™ XD Application Acceleration Solution, a product that caches storage area network (SAN) or direct attached storage (DAS) storage on server-side Nytro WarpDrive™ PCI-Express®-connected flash card. This evaluation was conducted in the Demartek lab in Arvada, Colorado and its purpose was to measure performance and latency improvements using a read-intensive decision support workload on an enterprise database software platform.

## Evaluation Summary

The LSI Nytro XD solution performed in an outstanding manner, increasing database application performance 11x over the baseline test. The Nytro XD solution also reduced latency by a factor of 10. All of these gains were obtained without changing the application or the back-end storage system.

We believe that the LSI Nytro XD solution approach is an excellent caching product for accelerating database performance with a server-side PCIe flash.

## Caching on Flash Technology

Caching on flash is garnering much attention for the promise of substantial performance increases and latency reductions for enterprise applications, especially database workloads. Caching on flash, regardless of where it is implemented in the I/O path, places a copy of "hot" data into its cache so that I/O activity can be accelerated. This caching technique benefits any application whose data is considered "hot" and is within the scope of management and visibility of the cache.

Flash-based caches improve performance over time as the cache fills up with hot data. This is known as "cache warm-up" or "cache ramp up" and can happen within minutes or hours, depending of the size of the cache and the I/O rate and access patterns.

### Server-Side Caching

While there are several possible ways to implement caching on flash, one excellent way is with the LSI Nytro XD solution composed of caching software and the Nytro WarpDrive™ PCIe flash-based card. The LSI Nytro XD solution fits into the category of "server-side caching." This type of flash-based caching accelerates I/O activity without requiring changes to the applications running in the host server. It also accelerates existing storage area network (SAN) storage without requiring any changes to SAN storage.

One of the primary benefits of server-side caching on flash is that as the cache inside the server warms, increasing numbers of I/O requests are satisfied without having to access the external storage network. This improves overall storage performance and drives down I/O response time, or latency. This is especially advantageous for financial and database applications that will benefit from lower latency.

## Performance Test Results

We ran a read-intensive decision support workload with the goal of measuring the increased performance offered by the LSI Nytro XD solution. We were especially interested in the amount of performance increase as the amount of data in the cache increased.

The database workload was run with the Nytro XD solution disabled, and then repeated with the Nytro XD solution enabled.

### Test Configuration

For these tests, we installed an enterprise database application in a Linux environment on a single server with an iSCSI storage device via a 10GbE connection.

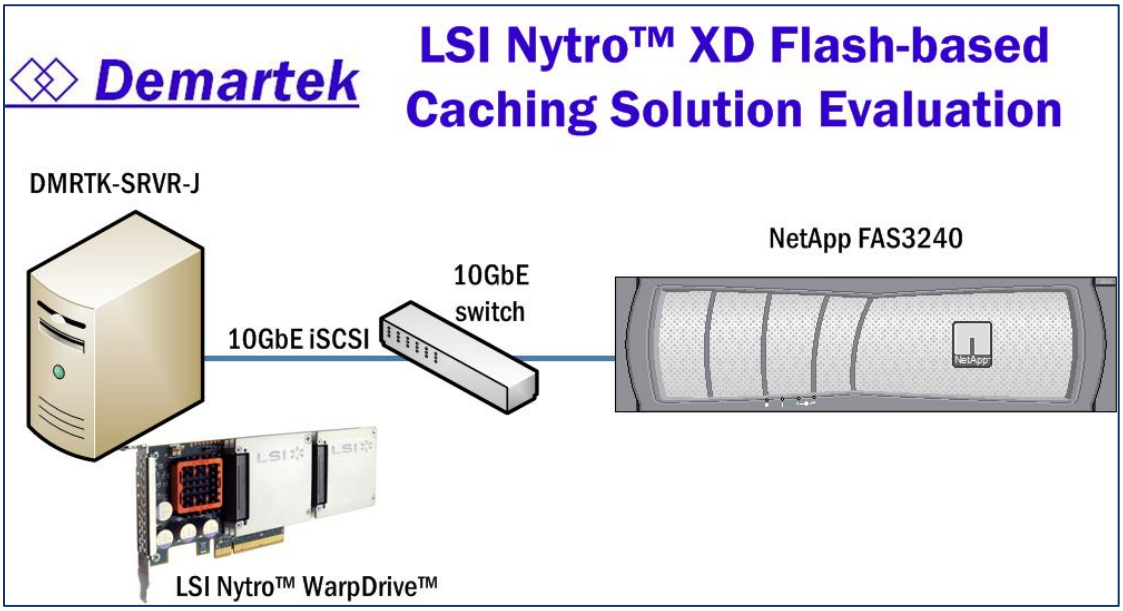

## Decision Support Workload

For the Decision Support workload, we repeated the test three times. The first time was with the cache disabled. Beginning with the second test, the cache was enabled. We repeated the tests with the cache enabled to see the effects of a pre-warmed cache, which would be the typical operating mode for a decision support workload where many ad-hoc queries are run on a slowly changing dataset.

This particular workload, known as TPC-H, consists of 22 specific sets of complex business-oriented ad hoc queries that examine a large amount of data. Each set of queries places a different load on the system. The test runs until all 22 sets of queries have been completed. When the full test suite is repeated with the same database and same parameters, the same amount of work is performed for each test suite.

The TPC-H workload performs significant sequential disk I/O as the single user queries perform index and table scans over the substantial datasets. The ability of storage to respond quickly allows the workload to complete faster. Lowering latency and increasing throughput are factors that enable quick responses.

### Time to Complete – Total Workload

These tests measure the time to complete the workload, so a lower score is better. As expected, the third test shows higher performance (completes in less time) than the first test with no cache and the second test with a partially warmed cache. The fully warmed cache performance in this test sequence was ***more than 11x faster*** than the baseline.

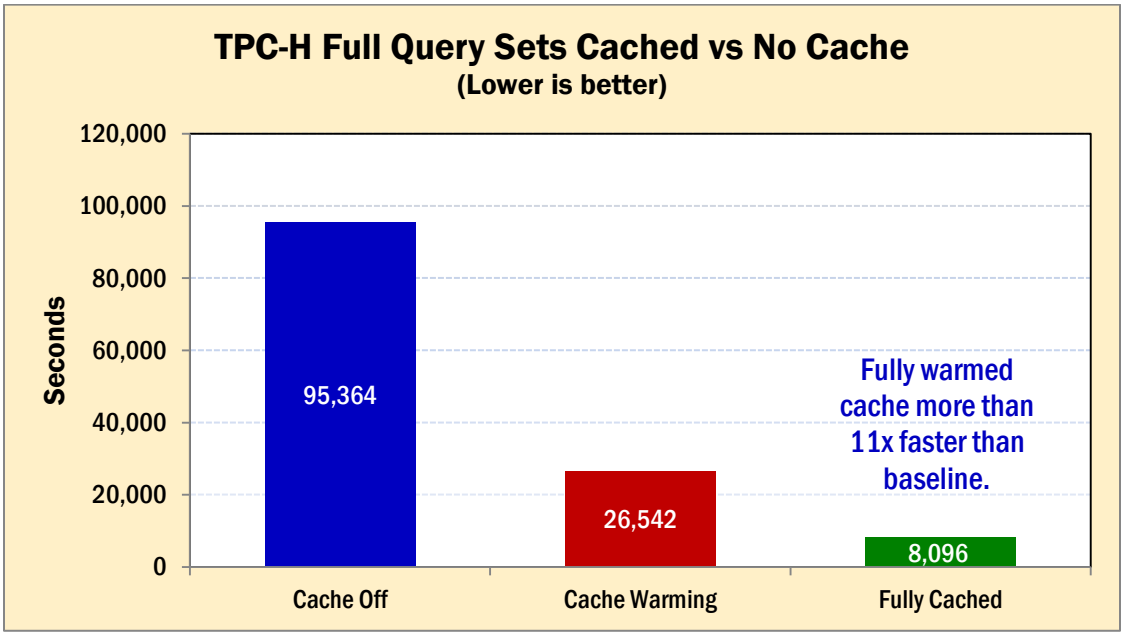

## Time to Complete – Individual Queries

This particular workload issues 22 different queries during the test. We captured the time to complete the workload by query for each of the three runs of the test. By the time of query 10 of run #2, the cache was warmed so the time to complete for the rest of run #2 and run #3 took full advantage of the flash-based cache.

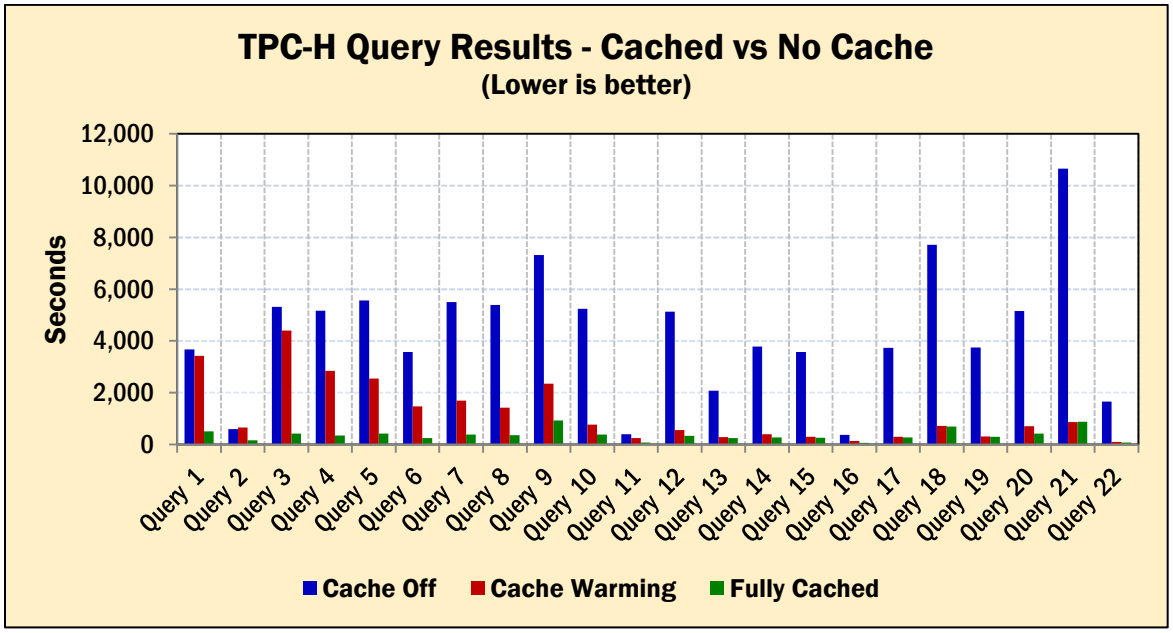

Latency

We measured the latency of these read-intensive tests, with and without the cache enabled. The data shows the obvious decline in the latency while the cache is warming (red line), with completion of the tests far faster than the baseline (blue line). The green line shows the fully warmed cache with the following latencies:

- Peak latency less than 0.5 milliseconds
- Average latency less than 0.2 milliseconds.

Completion of the full test suite, as noted previously, was more than 11x faster than the baseline. This test is based on the amount of work completed, so a faster storage system yields reduced completion times for the same amount of work.

An increasing number of applications are sensitive to latency. For example, in an OLTP application, short response times are critical. A transaction may require many successive queries to a database, where each query depends on answers returned from the previous query. User response times are entirely dependent on how quickly storage can return answers to the queries.

Busy web servers can also take advantage of reduced latencies. Two critical measurements for web servers are average page response time and average time to first byte (TTFB), both measures of latencies. SSD caching solutions such as the LSI Nytro XD solution are excellent tools to significantly reduce web server latency and improve the user experience.

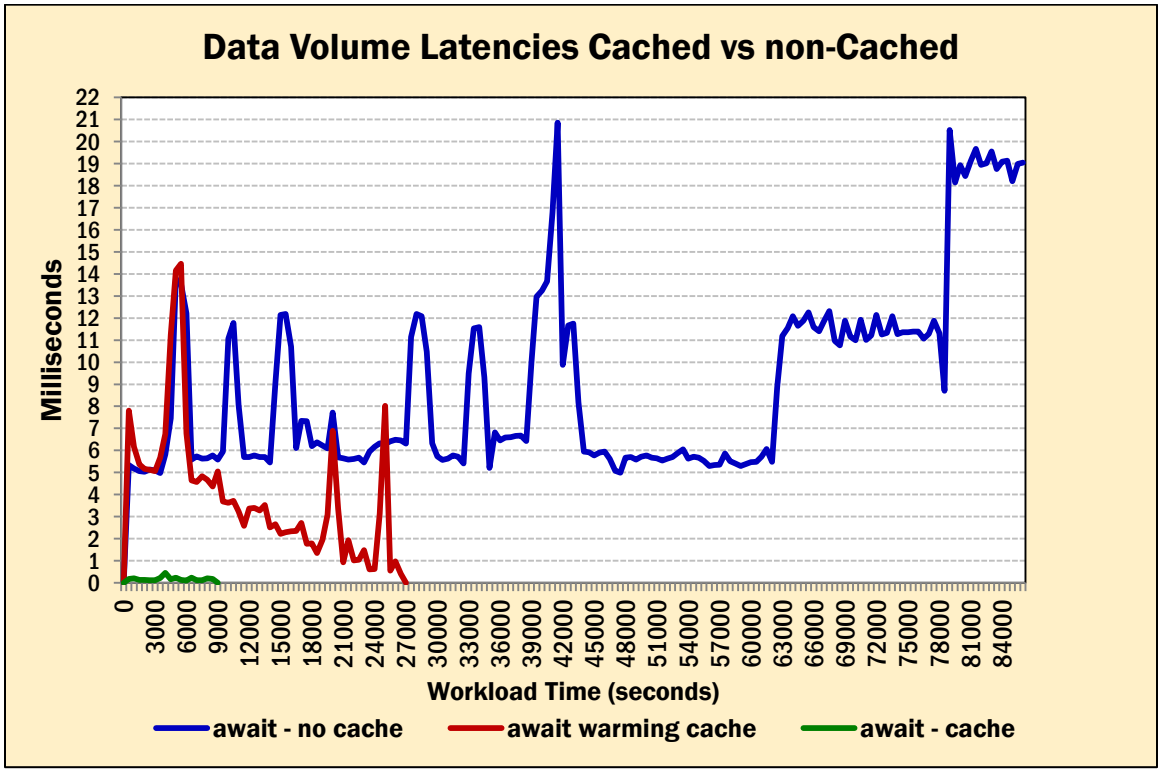

## Throughput (MBPS)

Throughput is the rate of data transfer between the source and the target. In this case, throughput means the amount of data read and processed by the database, in megabytes per second. For a fixed amount of work, higher throughput rates result in faster job completion.

We measured the throughput rate of these read-intensive tests, with and without the cache enabled. The LSI Nytro XD solution significantly accelerated the processing rate for this test. The peak and average throughput rates were:

- Peak throughput was 10x faster than the peak throughput of the baseline
- Average throughput was 11x faster than the baseline, once the cache was fully warmed

Note that the throughput achieved while the cache was warming (red line on the graph) started at the baseline rate and rapidly increased.

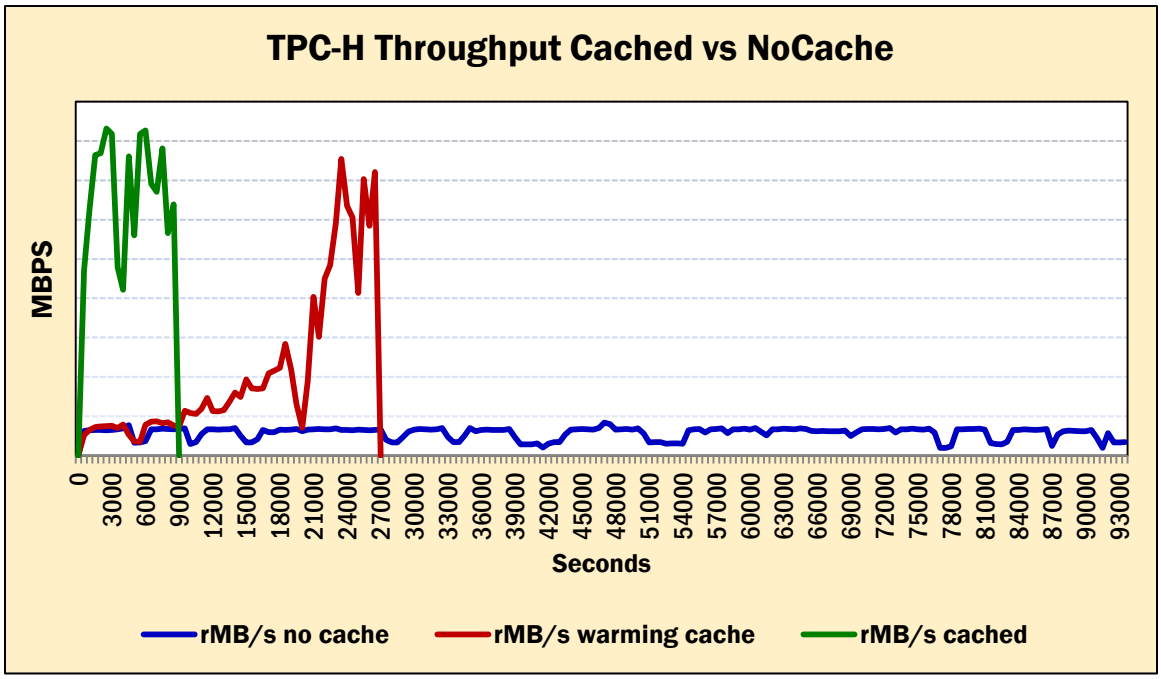

**Storage System Functions**

One of the important features of the LSI Nytro XD solution is that is does not adversely affect the normal operations of the back-end storage. In order to demonstrate this, we took a snapshot of the database while the Nytro XD solution's cache was enabled. We followed normal best practices by using the steps outlined below.

1. Enabled Nytro XD solution cache
2. Ran a database workload
3. Quiesced the database
4. Issued a snapshot command from the back-end storage system
5. Restored the data volumes from the snapshot
6. Brought the database back online
7. Ran another database workload

We were able to perform a database quiesce, a back-end storage system snapshot and other related functions while the Nytro XD solution's cache was enabled and all of the database and storage system functions operated correctly. As expected, the Nytro XD solution did not interfere with any of the storage system features and functions.

## Conclusion

The basic premise of an SSD cache is that the caching solution, in this case the LSI Nytro XD solution, identifies frequently accessed data, commonly known as "hot" data. LSI Nytro XD solution automatically moves ***a copy*** of this hot data to the SSD cache. This gives the following benefits:

- Any application that has "hot" data will benefit from the LSI Nytro XD solution, and this can be more than one application.
- Performance improves over time as the cache is populated with data, known as the cache "warm-up" period.
- The LSI Nytro XD solution reacts immediately to any data that becomes hot.
- The LSI Nytro XD solution does not require any changes to the application or the back-end storage configuration.
- The LSI Nytro XD solution does not require any data movement to different tiers of storage.
- The LSI Nytro XD solution is simple and easy to manage.

The LSI Nytro XD Application Acceleration solution performed in an outstanding manner, increasing database application performance by 11x when the cache became fully warmed. The Nytro XD solution also reduced latency by more than 10x to an average of less than 0.2 milliseconds. All of these gains were obtained without changing the application or the back-end storage system.

For those IT professionals seeking significant performance improvements, the Nytro XD solution is an excellent product.

Additional information regarding the Nytro XD solution is available at www.TheSmarterWayToFaster.com and from www.LSI.com/acceleration.

## Appendix – Evaluation Environment

All the tests were run in the Demartek lab in Colorado, and were run on a single server.

**Server – DMRTK-SRVR-J**

- 2x Intel Xeon X5680, 3.33 GHz, 12 cores, 24 threads ("Westmere")
- 144 GB RAM
- PCI-Express 2.0 slots
- RedHat Enterprise Linux 6.1, kernel 2.6.32-131.0.15.el6.x86_64
- Boot drive: SSD connected to a motherboard SATA port
- Local data storage: 2x 1 TB Seagate Barracuda ES.2 7200 RPM
- LSI Nytro WarpDrive 400 GB PCI-Express flash-based card
- Intel X520-SR 10GbE dual-port NIC
- Enterprise Database application

**Network Switch**

- Cisco Nexus 5020 10GbE switch

**Storage**

- NetApp FAS3240
- 2 TB local data storage in the server
- LSI Nytro WarpDrive 400 GB PCIe SSD used as an SSD cache

**Database Storage**

- 600 GB LUN on the FAS3240 allocated to the database
- Approximately 370 GB of data and indexes used by the database application
- 300 GB Redo logs allocated to local server data storage

---

The original version of this report is available at http://www.demartek.com/Demartek_LSI_NytroXD_Evaluation_2012-09.html on the Demartek web site.