



Improving Database Performance by Using the LSI® Nytro™ MegaRAID® Application Acceleration Card for DAS

Introduction

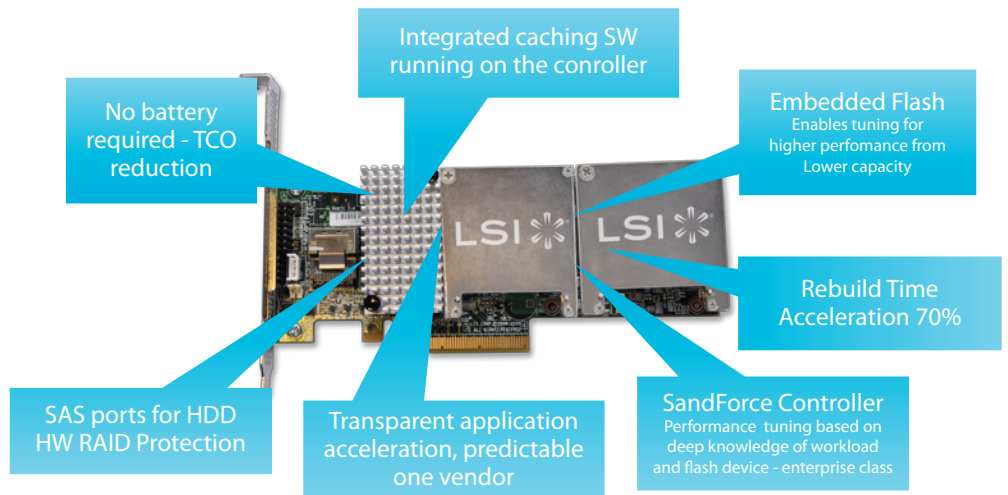
Supporting high transaction rates with minimal response times is frequently a requirement for today's enterprise databases, but many Database Administrators (DBA) don't have the time to successfully tune the database to meet these requirements. In some cases, companies will hire consultants to tune their databases to meet these application response time requirements. While tuning may help, often just improving the storage infrastructure can provide significant performance gains. A great storage solution to accelerate a database is to implement the LSI® Nytro™ MegaRAID® card. This paper will discuss the significant performance improvements that may be achieved with an Oracle® database application with little or no DBA involvement.

Latency Improvements to HDD

While server and network performance has managed to keep pace with the aggressive user workloads and the requirement for minimum response times, performance with storage-based hard disk drives (HDDs) hasn't improved appreciably. Transaction response times are often limited by high HDD latencies. This limitation has attracted much interest in using flash technology to supplement or replace HDDs in certain applications. In many cases it would be cost prohibitive to move to a 100% flash solution. One, more economical approach, is to use a small amount of low latency flash for frequently accessed data and HDDs for the rest of the system needs.

Application performance can increase significantly by reducing the latency to frequently accessed data, known as "hot spots". These hot spots frequently occur in databases due to frequently accessed data and indexes. Caching these hot spots on flash storage can help provide high performance with a low investment cost compared to replacing the DAS devices.

The Nytro MegaRAID card embodies LSI's approach to not only make fast solid-state storage (SSD) products, but to make them safe and easy to use. The Nytro MegaRAID card combines flash for performance and HDDs for economy. It integrates RAID protection, caching intelligence, and embedded flash on a single PCIe® card. It utilizes the LSI SandForce® flash controller family that has shipped millions of units and a RAID controller that has been vetted by millions of MegaRAID controller installations.



Before the Nytro MegaRAID card was an option, the use of flash to accelerate a database was more difficult to implement and the DBA would have to consider the following options to increase performance:

- Placing all of the data on flash.
- Manually persisting data between disks and SSDs.
- Implement a SAN that supports SSD and disk tiers.

Placing everything on flash is an expensive proposition and is typically not necessary. A general rule of thumb for Online Transaction Processing (OLTP) environments is that 20% of the database receives 75% of the storage access. This formula validates the usefulness of caching database data on a small amount of flash storage.

Manually placing data between HDDs and flash can be a time-consuming management nightmare. Storage systems are typically not configured to select just the “hot” data blocks, but require moving larger storage segments which results in wasted flash capacity. Additionally, changes over time require continual tweaks to the data placement and additional DBA effort is required to ensure that the data protection schemes between the disk and flash are synchronized.

Using a Storage Area Network (SAN) that supports SSD and HDD tiers is an option for SAN environments, but it comes with significant drawbacks. Instead, moving SAN-attached flash to the PCIe bus can help increase performance and decrease costs.

The Nytro MegaRAID card utilizes flash and HDDs to provide both capacity and performance improvements. The card utilizes LSI enterprise acceleration and data protection technology. It is designed not for “performance at any cost,” but as a solution to bring flash performance to mainstream applications.

Goal

To benchmark the performance of an Oracle database configured for an OLTP application using the LSI Nytro MegaRAID card with a DAS-based HDD compared to the same setup with traditional DAS-based HDD storage.

System Infrastructure

For these tests, one server internally in an LSI lab was running an Oracle 11gR2 database and the tests were replicated. The database was restored between each benchmark. The first test was to establish a baseline using all DAS-based HDD. The second test was executed after implementing the Nytro MegaRAID card to cache the virtual disks with the database data and index components.

Quest® Benchmark Factory® software for databases was used in the above testing scenario. Benchmark Factory software is a database performance testing tool that allows you to conduct database workload replay, industry-standard benchmark testing, and scalability testing. Quest, Quest Software, and Benchmark Factory are registered trademarks of Quest Software, Inc. in the U.S.A. and/or other countries. All other trademarks and registered trademarks are property of their respective owners.

NOTE: None of the results from LSI's in-house testing were audited or published by the Transaction Performance Council (TPC), the copyright owner of TPC-C. The TPC-C transactions are defined according to the TPC-C standard specification.

Hardware Specifics

- 2 CPUs, Quad Core, Intel® E5640 @ 2.67GHz
- 48GB memory
- LSI Nytro MegaRAID Card
- 20 – 73GB SAS HDD
- RAID 1 – all HDD in Write Back mode
- flash cache to hold the hot data blocks (20% of the database)

Operating System Specifics

- Red Hat® Enterprise Linux® 6.0
- Kernel Linux 2.6.32-71.el6.x86_64

Database Specifics

- Oracle 11gR2 Enterprise Edition
- 2.8GB SGA
- Use Filesystems
- Data and Indexes on the cached virtual disks
- Asynchronous IO
- 8k Block Size
- 8 - DBWRs

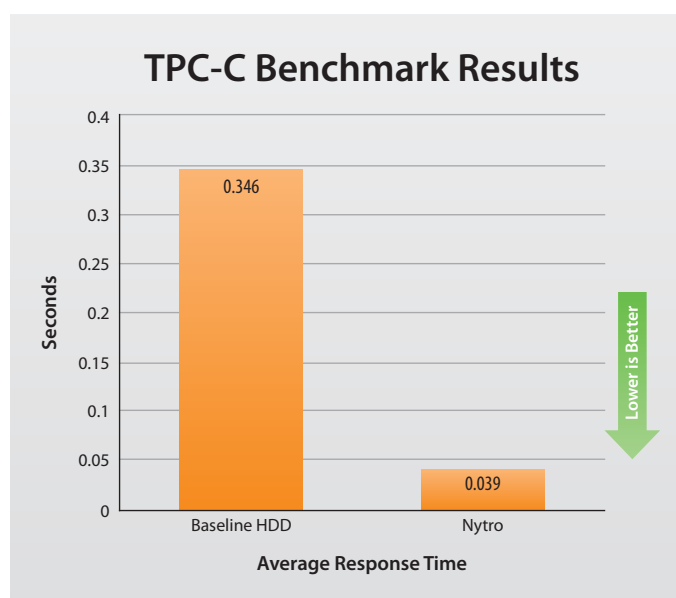
Benchmark Factory Software Specifics

- Benchmark Factory 6.61 software
- Benchmark Factory TPC-C software Benchmark settings:
 - Scale - 900
 - Users - 200
 - Interarrival Time – Absolute – 500ms
(Used Interarrival Time to stress the storage and database)

Transaction Results

A goal to database tuning is to decrease overall transaction response times within the database. If database transactions finish faster, the overall application performance is accelerated. Implementing the LSI Nytro MegaRAID card is a method of tuning the database, but without any DBA involvement.

The benchmark results showed substantial performance gains in the average response times by utilizing the Nytro MegaRAID card to cache DAS-based HDD data. When running the benchmark to establish a baseline using all DAS-based HDD, the average response time of 0.346 of a second was achieved. Using the Nytro MegaRAID card for caching, performance jumped almost 9x over the DAS-based HDD baseline.



Summary

From the multiple internal LSI tests and benchmarks that were executed as described herein using an Oracle database simulating a real-world online transaction application, a 9x reduction in latency was realized by implementing the Nytro MegaRAID card with a DAS-based HDD array. The benchmark results support that the implementation of the Nytro MegaRAID card can have the following benefits in a DAS environment:

- Substantial performance gain over certain HDD environments
- A cost effective way to increase performance
- Lower cost solution compared to short stroking hundreds of HDDs to get the equivalent IOPs from the Nytro MegaRAID SSD card
- Can be used with any Oracle database
- Designed for easy integration into an existing DAS storage environment, without involvement from the database or application administrator

LSI Corporation makes no representations, warranties or guarantees regarding the compatibility or performance of, or results generated by, specific hardware or software products. The information herein is provided "AS IS" without warranty of any kind express or implied for use at your sole risk. Without limiting the foregoing, the warranties of merchantability and fitness for a particular purpose are expressly disclaimed. This document is not warranted to be error-free. LSI Corporation has no responsibility to update the information herein and reserves the right to make changes to this document and to any of its products and services at any time without notice. LSI Corporation does not assume any responsibility or liability arising out of any reliance on or use of this information. Individual results may vary.

For more information, visit: www.lsi.com/acceleration

1. S. T. Leutenegger and D. M. Dias. A Modeling Study of the TPC-C Benchmark. In SIGMOD, 1993.

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



North American Headquarters
Milpitas, 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
Tel: [+81] 3.5463.7165

LSI, LSI & Design logo, MegaRAID, and Nytro are trademarks or registered trademarks of LSI Corporation. All other brand or product names may be trademarks or registered trademarks of their respective companies.

LSI Corporation reserves the right to make changes to any products and services herein at any time without notice. LSI does not assume any responsibility or liability arising out of the application or use of any product or service described herein, except as expressly agreed to in writing by LSI; nor does the purchase, lease, or use of a product or service from LSI convey a license under any patent rights, copyrights, trademark rights, or any other of the intellectual property rights of LSI or of third parties.

Copyright ©2012 by LSI Corporation. All rights reserved. > 0312