

Increase Database Performance by Implementing Cirrus Data Solutions' DCS SAN Caching Appliance With the Seagate Nytro Flash Accelerator Card

Technology Paper

Authored by Rick Stehno, Principal Database Engineer, Seagate

Introduction

Supporting high transaction rates with minimal response times is frequently a requirement for today's enterprise database customers. But many Database Administrators (DBAs) don't have the time to successfully tune databases to meet these demands. In some cases, companies need to spend a significant amount of time and money to hire consultants to tune their databases in order to meet the required application response time. While tuning may help to a certain extent, often the real problem is at the storage level, and just improving and updating the storage infrastructure can provide significant performance gains. One example of such a Storage Area Network (SAN) solution for accelerating databases is Cirrus Data Solution's Data Caching Server (DCS) Appliance. This paper will discuss how DCS delivers significant performance improvements to an Oracle database application with no DBA involvement, no changes to existing hosts, fibre channel (FC) fabrics, FC storage, and with no downtime.

Latency Improvements to a SAN

While server and network performance keeps pace with the aggressive user workloads and the requirement for minimum response times, performance with traditional storage implementations hasn't improved appreciably. This limitation attracts interest in using flash technology to supplement or replace existing storage deployments. However, in many cases it is cost prohibitive to abandon existing investments of traditional storage and move to a 100% flash solution. A more economical approach is to use smaller amounts of low latency flash for frequently accessed data and traditional storage for the bulk of data that is not frequently accessed.

Application performance can increase significantly by reducing the latency to accessed data, known as *hot spots*. These hot spots often occur in databases due to frequently accessed data and indexes. Caching these hot spots on flash storage can provide higher performance with a low investment cost compared to replacing the entire SAN-based storage farm.

Increase Database Performance by Implementing Cirrus Data Solutions' DCS SAN Caching Appliance With the Seagate Nytro Flash Accelerator Card



Cirrus Data Solutions' Data Caching Server Appliance

Cirrus Data Solutions' Data Caching Server (DCS) Appliance is a 2U Linux-based server equipped with multiple pairs of FC ports and with up to three Seagate Nytro flash accelerator cards (and another ten Nytro cards when used with an optional external PCIE extended enclosure) and 512GB of RAM. Always deployed as a highly available (HA) cluster pair, the DCS can support up to 44.8TB of PCIe-based flash card capacity and 1.0TB of DRAM as cache resources for acceleration of the SAN based database storage. The DCS appliances can be inserted live between the hosts and the FC SAN storage without requiring any changes to the application hosts, FC switches (zoning) or the FC storage controllers (LUN masking). Simply physically reconnect the storage controller ports, one port at a time, through the FC ports of the DCS appliance, and it is ready to be put into action. The DCS automatically discovers all the hosts, storage controllers and LUNs. The administrator can then configure the appliance to cache data on the SAN by enabling Cache Policies for various groups of LUNs. Since not all of the SAN data is hot, the DCS will cache only the data blocks that are frequently accessed on the LUNs in which the DBA specifies. In addition, DCS can also dynamically configure the cache based on discovered hot spots. This cached data will reside on the appliance's PCIe-based flash cards; such as the Seagate Nytro flash accelerator cards.

The DCS solution is designed to be dynamically inserted between the host and the SAN with no interruption to service. The data cached on the appliances is done using Write-Through technology in which the data is always committed on the SAN storage in addition to being stored in the cache. Cached data stored on the appliances is protected by the permanent storage on the SAN-based (hard disk drive) HDD volumes. For reads, the requested block is read from the cache storage (flash or DRAM) if already cached. If not, the requested block is read from the SAN-based HDD into the cached storage, making it available for future access. Over time, the cache storage is always filled with the hot data of that particular moment of time, making it possible for a small amount of cache storage to accelerate the overall traditional SAN-based HDD volumes. All this can be accomplished with little or no involvement from the DBA, nor will there be any need for changes to the enterprise database.

DCS is built on a patent called Transparent Data Intercept (TDI). This allows for Zero Change, plug-and-play development of DCS in the SAN where the paths are inserted one at a time,

five seconds per path in a live environment. The cache can be a combination of DRAM, SSDs and/or PCIe flash cards. DCS is built on an Evolution engine that utilizes statistical data to help provision the cache.

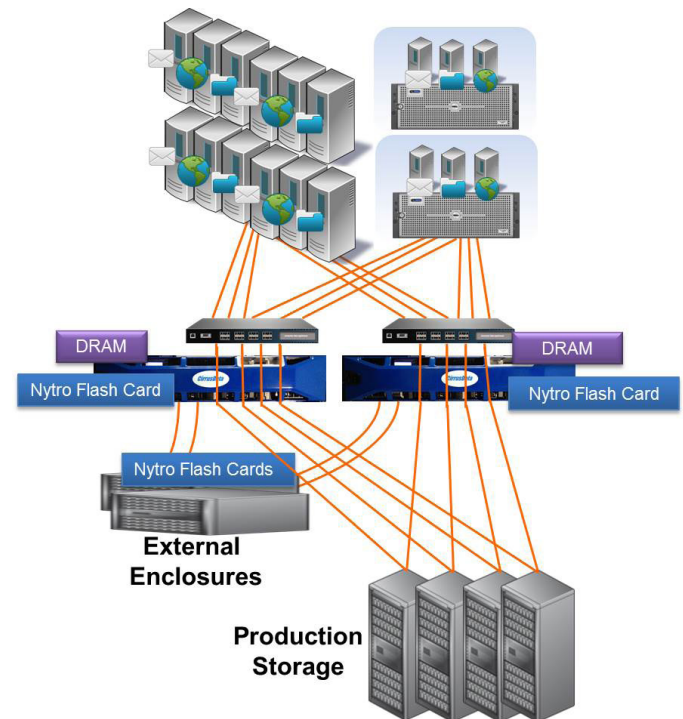


Figure 1: CDS DCS: Transparent Data Intercept

Before the DCS solution was an option, using 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 across both HDDs and flash
- Implementing a SAN that supports SSDs and HDDs and automated tiering.

Placing everything on flash is expensive 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 access¹. This formula validates the usefulness of caching database data on a small amount of flash storage.

Manually persisting data on SAN-based HDDs and flash can be a time-consuming management nightmare. Database files are typically a few large files, with hot and cold data all mixed

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

Increase Database Performance by Implementing Cirrus Data Solutions' DCS SAN Caching Appliance With the Seagate Nytro Flash Accelerator Card



within. It is often impossible to manually create files to store just the hot data blocks. Instead, the entire database file will need to be moved to flash, resulting in wasted flash capacity. Additionally, changes over time require continual updates to the data placement, making it impossible to keep up.

In SAN storage environments valuable data management services may be running on the storage controller, such as disk-to-disk backups, replication and snapshots. Manually moving some of this data onto a separate SSD resource outside of the SANs management framework is not an option without changing the data management methods. This is a major hurdle that must be overcome in the way of taking advantage of the performance of flash solutions that are available directly on a server's PCIe bus.

The use of caching on servers for accelerating access has been around for decades. The DCS appliances allow the administrator to choose which data on the SAN needs to be cached for best performance. The appliance directly copies frequently accessed data blocks to the cache to help deliver low latencies and greatly improves performance more than SAN-based HDDs can achieve. As hot data *cools*, the DCS Cache Policy tracks data temperature and automatically caches the hotter data and not the cooler data. Operating in write-through mode, any storage services provided by the SAN controllers are not disrupted.

Nytro Flash Accelerator Card

The Seagate Nytro flash accelerator card is a solid-state primary storage solution for accelerating database workloads. The small PCIe footprint gives database administrators the option of easily transforming their datacenter infrastructure to flash. Built with advanced solid-state capabilities and technology that optimizes endurance and reliability, the Seagate Nytro flash accelerator card offers an enterprise ready solution.



Figure 2: CDS DCS: Nytro Flash Accelerator Card

CPU resources are also under a great deal of pressure to process the data in today's demanding environments. The CPU can become a performance barrier if its processing resources are pulled away to other tasks. Nytro cards, in comparison to other flash technologies, utilize host-offload architecture for I/O and flash management, which results in consistent performance and low host CPU utilization to enable the CPU to be fully available for its primary task of processing data.

Demonstrated Testing Goal

The goal of the following tests were to benchmark the performance of an Oracle database configured for an OLTP application using Cirrus Data Solutions' DCS Appliance with the Nytro flash accelerator card and compare this to a traditional SAN-based configuration.

System Infrastructure Setup

For these tests, a Linux server in the Cirrus Data Solutions lab was running an Oracle 11gR2 database, and a Windows server was running a Quest Benchmark Factory OLTP benchmark. Since it is easy to deploy the DCS Appliances, the first set of tests was done against the SAN-based HDD to establish a baseline for the Average Response Time and Transactions per Second (TPS) statistics. The second set of tests was done by inserting the DCS Appliances, with the Nytro cards, between the database server and the SAN-based configuration. The database was restored before running the benchmark. The SAN-based configuration's database, with the DCS Appliances, was configured to cache only the Logical Units (LUNs) with the data and indexes components.

Quest Benchmark Factory software for databases was used in the above testing scenario and is a database performance-testing tool that allows you to conduct database workload replay, industry-standard benchmark testing and scalability testing.

NOTE: None of the results from this 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.

Increase Database Performance by Implementing Cirrus Data Solutions' DCS SAN Caching Appliance With the Seagate Nytro Flash Accelerator Card

Cirrus Data Solutions DCS Appliance

- DCS-2000 with 1.6TB Nytro flash accelerator cards
 - DCS can accommodate up to three Nytro cards (3 × 3.2TB = 9.6TB) per appliance, which is 19.2TB per pair of a DCS cluster
 - DCS with an optional external PCIe extension enclosure can have up to 14 Nytro cards (14 × 3.2TB = 44.8TB)
 - Provisioned 50GB Cache for Oracle

Hardware Specifics for Database Server

- Dell R720 with two CPUs, six Cores, Intel Xeon CPU E5-2630 @ 2.30GHz
- 24GB memory
- Two 4Gb FC Qlogic ports

Storage

- CX-700 with 29 × 60GB SAS HDD, no RAID. Total usable data is 1.74TB
- Oracle ASM

Operating System Specifics

- Oracle Linux Server release 6.5
- Kernel Linux 3.8.13-26.2.1.el6uek.x86_64

Database Specifics

- Oracle 11gR2 Enterprise Edition
- Asynchronous I/O
- 8k Block Size

Benchmark Factory Software Specifics

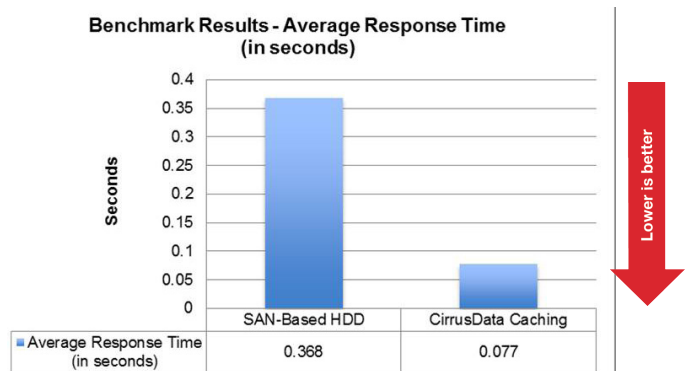
- Benchmark Factory software 6.61
- Benchmark Factory TPC-C software benchmark settings
 - Scale - 800
 - Users - 100

Expected Transaction Results

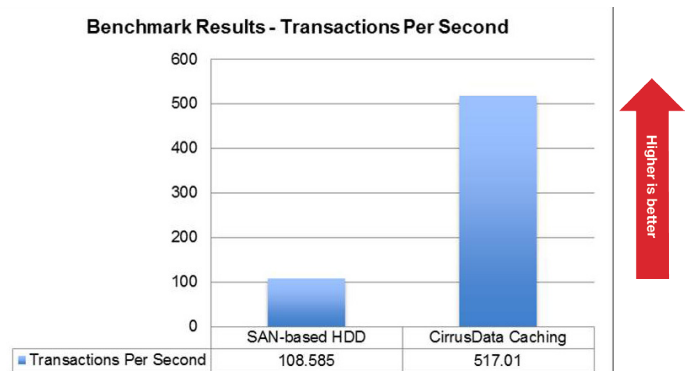
One main goal of database tuning is decreasing overall transaction response times within the database. If database transactions finish faster, the overall application is accelerated. Implementing DCS in a SAN-based configuration is a method of tuning the database, without any DBA involvement or disruption to the services provided by the SAN.

Benchmark Results

The benchmark results showed substantial performance gains in the Average Response Times and in the TPS areas by utilizing Cirrus Data Solutions' DCS Appliances with the Nytro flash accelerator card to cache SAN-based HDD data. When running the benchmark to establish a baseline using all SAN-based HDD, the average response time of .368 seconds was achieved. Using the DCS Appliances for caching, performance jumped almost 5x over the SAN-based HDD baseline, resulting in an average response time of .077 of a second.



When running the benchmark to establish a baseline using all SAN-based HDD, TPS of 108.585 was achieved. Using the DCS appliances for caching, performance jumped almost 5x over the SAN-based HDD baseline, resulting in a TPS result of 517.01.



Increase Database Performance by Implementing Cirrus Data Solutions' DCS SAN Caching Appliance With the Seagate Nytro Flash Accelerator Card



Summary

From the multiple benchmarks that were executed using an Oracle database simulating a real-world online transaction application, a 5× reduction in latency was realized by implementing the DCS Appliance solution and the Nytro flash accelerator cards with a SAN-based HDD array. The benchmark results show the implementation of the DCS solution can have the following benefits in a SAN environment:

- Substantial performance gain over SAN-based configuration, making it well-suited to be used complementary to a SAN
- A more cost-effective way to increase performance
- Easy integration into an existing SAN storage environment since as it requires no involvement from the database or application administrator to implement

NOTE: Seagate successfully completed all of these system and storage modifications and benchmarks in the Cirrus Data Solutions internal performance lab. Before implementing any of these modifications in your specific environment, be sure to test them completely to determine whether to use them in your environment.

seagate.com

AMERICAS Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000
ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16-18, rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00

© 2015 Seagate Technology LLC. All rights reserved. Printed in USA. Seagate, Seagate Technology and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Nytro is either a trademark or registered trademark of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. Cirrus Data Solutions and the Cirrus Data logo are trademarks or registered trademarks of Cirrus Data Solutions Inc. All other trademarks or registered trademarks are the property of their respective owners. When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Actual data rates may vary depending on operating environment and other factors. Seagate reserves the right to change, without notice, product offerings or specifications. TP668.1-1501US, January 2015