



Reduce Latency and Increase Application Performance Up to 13x With Adaptec maxCache™ 2.0 Read and Write SSD Caching Solutions

Executive Summary

Today's data centers and cloud computing environments require increased I/O performance and decreased latency in order to support large-scale applications such as web serving, file serving, databases, online transaction processing (OLTP), Microsoft Exchange Server, and high performance computing (HPC). This whitepaper focuses on the performance benefits of using Adaptec maxCache 2.0 Solid State Drive (SSD) Caching Solutions for these key data center applications and the cost savings that can be realized from this improved application performance.

The performance and financial benefits of maxCache 2.0 are illustrated below by comparing the performance levels in a 100% Random-Read IOMeter workload scenario with a maxCache 2.0-enabled Adaptec storage controller vs. an equivalent Adaptec storage controller without maxCache 2.0, as well as an OLTP scenario with a maxCache 2.0-enabled Adaptec storage controller vs. an Adaptec storage controller without maxCache 2.0. These quantified performance improvements can be applied to a Total Cost of Ownership (TCO) calculation summarizing the potential savings in Capital Expenses (CapEx) and Operating Expenses (OpEx).

Performance Highlights

Accelerating application performance, reducing latency, and increasing the number of users supported per server are key requirements for data centers and cloud computing environments. Tests of Adaptec maxCache 2.0 SSD Caching when compared to hard disk drive (HDD)-only solutions show:

- **Up to 13x performance improvement** in read-intensive I/O operations per second (IOPS)
- **Up to 13x latency reduction** in read-intensive applications

Comparing a SAS HDD-only solution to a maxCache 2.0-enabled SATA HDD/SSD solution, observations were:

- **6x performance improvement** in mixed workload IOPS
- **10x latency reduction** in mixed workload applications
- **7x increase in available server capacity** and significantly-reduced cost per GB and cost per I/O

Introduction

Data centers and cloud computing environments require application-tuned, high-density servers to provide end-users with consistently high quality of service (QOS). Usually, these server

deployments are targeted towards specific applications such as web serving, file serving, databases, OLTP, Microsoft Exchange Server, and HPC. Customers of these data centers and cloud computing applications typically sign service level agreements (SLAs) that obligate the operators to provide pre-determined levels of reliability, availability, and serviceability (RAS) as well as performance. To meet these requirements, data center operators deploy servers with high I/O throughput and high IOPS that are "application-tuned" to ensure adherence to the SLA. Violating the SLA terms can lead to severe financial penalties for the data center operator, not to mention the adverse customer reaction to any unsuitable performance or user experience.

In addition to SLA requirements, budgets compel data centers to increase the number of hosted users per server in order to reduce the cost of service (COS) per user as well as the CapEx and OpEx of additional hardware.

Adding more end-users to a server will cause an increase in latency and a decrease in IOPS. Since both of these issues happen well before storage capacity is maxed out, data centers and cloud computing environments are forced to add new servers to maintain appropriate levels of service, even if the existing servers have the storage capacity for more users. This results in low storage capacity utilization and an increase in CapEx and OpEx (namely maintenance, power and cooling costs), as well as an increase in physical space requirements.

Additionally, to meet the fast response time and availability requirements of the SLA and increase the number of hosted users required for reducing the COS, operators tune servers to pool data in the system cache despite the fact that not all data is accessed by users uniformly. Web servers, for example, pool homepages of the websites they host much more frequently than other pages. E-commerce applications query images of frequently-accessed ("hot") products much more often compared to other ("cold") products, and also sell these "hot" products more often.

But it is not just the SLA that is at stake. Increasingly, end users and customers are becoming less tolerant of latency and poor IOPS performance. According to Equation Research:¹

- 78% of site visitors have gone to a competitor's site due to poor performance during peak times.
- 88% are less likely to return to a site after a poor user experience.
- 47% left the site with a less positive perception of the company.

A poor-performing website can have quantified financial implications. Amazon has found that a mere 0.1 seconds of latency shaves 1% from sales revenues, for example.¹

¹ "When more Website visitors hurt your business: Are you ready for peak traffic?," Equation Research, 2010

Adaptec maxCache 2.0

Companies that address the latency challenge have reaped rewards. Shopzilla, for example, reports that improving website performance by five seconds led to 25% more page views and a 12% revenue increase; large financial institutions leverage a 0.5-second speed advantage to execute millions of orders per second and make record profits.¹

Technology that can provide fast responses for “hot” data while maintaining a large capacity of “cold” data will result in direct savings in the number of servers deployed. In other words, any improvement in the servers’ I/O capability increases the number of users per server and lowers the cost per user. A technology that can deliver this improved performance while also increasing the available server capacity will improve the cost per user and user per server metrics even further.

Adaptec maxCache 2.0 SSD Caching meets the needs of I/O-intensive data center and cloud computing environments by allowing them to convert industry-standard servers into cost-effective, high-performance, scale-out application storage appliances that optimize critical usage, performance, and financial metrics.

About Adaptec maxCache SSD Caching Solutions

Adaptec maxCache combines SSDs used as cache with Adaptec’s maxCache SSD Caching software to dramatically accelerate I/O performance and reduce costs without disrupting existing operations. It virtually eliminates the bottleneck that can occur between CPUs, memory and storage.

SSDs offer a number of advantages compared to HDDs, including higher read bandwidth, higher IOPS, better mechanical reliability (due to the absence of moving parts), and higher resistance to shocks and vibrations. However, the same features (i.e., flash) that provide these advantages also come with some inherent limitations compared to HDDs, such as limited capacity and lower streaming write bandwidth. An SSD’s lifespan is also highly dependent on the number of write operations performed to the device.

To capitalize on the advantages of SSDs while suppressing their limitations, maxCache SSD Caching delivers performance benefits for both read and write operations. A patent-pending Learned-Path Algorithm identifies frequently-accessed data and optimizes reads and writes by moving a copy of this data directly into an SSD cache for faster retrieval of future requests. Storing only the “hot” data in the SSD cache optimizes the balance of SSD performance and capacity. By leveraging its unique presence in the data path to create a “cache pool” of “hot” data, maxCache 2.0 SSD caching can provide significant performance gains compared to HDD-only deployments (Figure 1). In read-only workloads, maxCache 2.0 delivers up to a 13x improvement in IOPs and a 13x reduction in application latency.

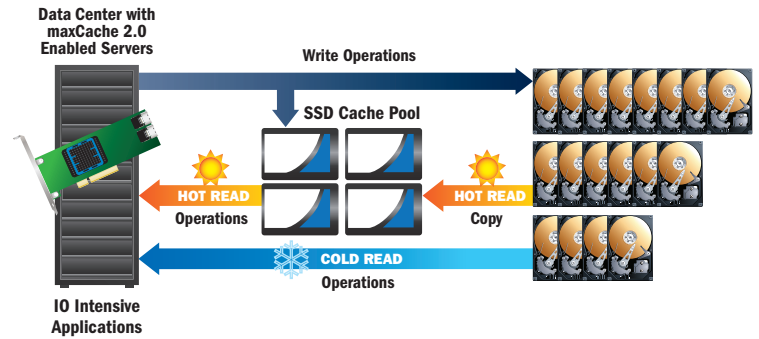


Figure 1. maxCache 2.0 Deployment

Adaptec maxCache 2.0 also enables caching of write data to bring the advanced performance and latency capabilities of SSD caching to workloads that mix reads and writes. By expanding the use of SSD caching, maxCache 2.0 is suitable for wide scale data center deployments, offering increased financial benefit while streamlining hardware design and implementation.

Optimization of available space is critical to current and future data centers considering the direct cost to design and wasted investment if space is unused. With the acceleration of IOPs and reduction in latency, maxCache 2.0 allows data center and cloud computing environments to host more users and perform more transactions per second while reducing the overall number of servers required to fulfill any given workload. This reduces a company’s CapEx significantly, leading to improved utilization of hardware as well as a reduced server footprint requirements within a data center.

The reduction in servers has an additional financial benefit of reducing the associated OpEx of power, cooling, and maintenance, delivering a highly-reduced TCO solution (Figure 2). OpEx reductions are a continued savings throughout the server lifespan, so the longer the server is in use, the greater the accumulated OpEx will be.

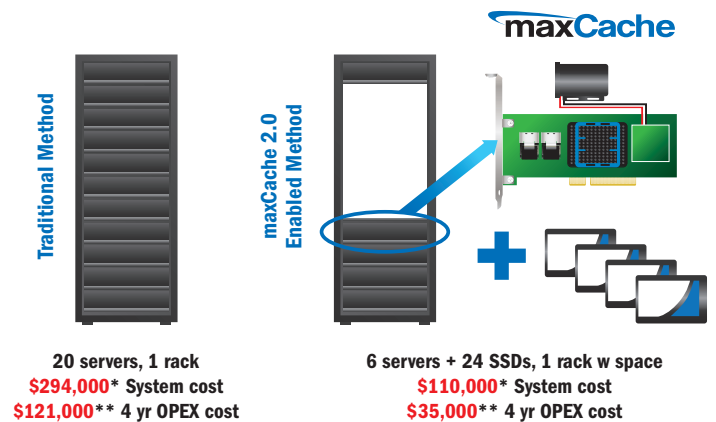


Figure 2. Quantified Impact on Utilization of Hardware Investment

*Using an estimated unit H/W cost of \$15,000 and zero S/W license fees

**Using system Power Estimator @ 800W per server, and a National DOE average of \$0.11 per kWh

¹ “When more Website visitors hurt your business: Are you ready for peak traffic?,” Equation Research, 2010

Adaptec maxCache 2.0

Adaptec maxCache 2.0 vs. standard SSDs and PCIe-based flash cards

Flash-based storage devices are becoming more affordable and prevalent in computing environments, but they should not be confused with maxCache 2.0 solutions. In order to get maximum performance when using a standard SSD or a PCIe-based flash card, applications must be tuned to store data that requires higher IOPs on the high-performance SSD or flash. This requires that an administrator have intimate knowledge of that specific application and understand how it will benefit by routing data that requires higher IOPs to the SSD.

Additionally, standard SSD storage capacities are still relatively small compared to those of rotating disk drives. Therefore, a larger number of SSDs are needed, which will increase CapEx significantly. With Adaptec maxCache 2.0 SSD Caching, one can still benefit from the larger capacities of rotating media (HDDs), while gaining the improved I/O performance benefits of SSDs. maxCache 2.0 delivers an optimized balance of SSD performance and HDD capacity to solve the unique challenges faced by data centers and cloud computing operators.


Test Methodology

IOmeter was used to compare performance of Adaptec maxCache 2.0 SSD Caching Solutions using an Adaptec Series 6 and Adaptec Series 6Q Storage Controller.

Tests were run to measure the benefits of maxCache 2.0 along two axes. The first set of tests was run to benchmark performance of maxCache 2.0 in a 100% read-only workload and compare against an HDD-only solution. A second set of tests was then run to determine the benefits read and write caching could deliver using high-capacity SATA HDDs and comparing against and HDD-only configuration of SAS drives.

Read Caching Performance – IOPs and Latency

Read-intensive applications such as web servers, file servers, and e-commerce applications radically benefit from maxCache 2.0, seeing up to 13x improvement in IOPs, and up to 13x reduction in latency.



- RAID 0 performance comparison under 100% Random Read IOmeter workload.


- HDD-only configuration: Eight 15k RPM SAS HDDs, 300GB capacity each, in RAID 0.
- maxCache 2.0 configuration: Eight 15k RPM SAS HDDs, 300GB capacity each, in RAID 0; two 100GB SATA SSDs for maxCache cache pool.

Write Caching Performance

The introduction of write caching support extends maxCache 2.0 benefits to I/O-intensive applications with mixed workloads, including OLTP, Microsoft Exchange Server, and HPC environments.

Write Caching Performance – IOPS & MB/s

Enabling maxCache 2.0 delivered significant benefits in performance while also allowing for a server design with much higher capacity potential. Performance in this example yielded a 6x increase in IOPS and MB/s. These performance gains were achieved by eliminating a SAS HDD server configuration in favor of a maxCache 2.0 configuration with high-capacity SATA HDDs. The configuration change expanded the available solution capacity from 2.4TB to 18TB – a 7x increase that significantly reduces a data center's cost per GB and cost per I/O metrics.




- RAID 5 performance comparison under OLTP workload (2:1 ratio of read:write operations).
- HDD-only configuration: Eight 15k RPM SAS HDDs, 300GB capacity each, in RAID 5.
- maxCache 2.0 configuration: Six 3TB SATA HDDs for RAID 5; two 100GB SATA SSDs for maxCache pool.

Write Caching Performance – Latency

Tests with maxCache 2.0 enabled confirmed a 6x reduction in average latency and a 10x reduction in maximum latency, as well as significant improvements in response time and the consistency of response times. Again, these measures compare a SAS HDD-only configuration to a maxCache 2.0 configuration using high-capacity SATA HDDs.

Adaptec maxCache 2.0



- RAID 5 performance comparison under OLTP workload.
- HDD-only configuration: Eight 15k RPM SAS HDDs, 300GB capacity each, in RAID 5.
- maxCache 2.0 configuration: Six 3TB SATA HDDs for RAID 5; two 100GB SATA SSDs for maxCache pool.

Conclusion

Data center operators and cloud applications are continuously challenged to improve server performance to keep up with the demands of high-throughput applications and growing user bases. At the same time, space restrictions, power and cooling limitations require data centers to find the most cost-, space-, and energy-optimized products to enhance the server's I/O and workload capability.

As highlighted in the tests above, maxCache 2.0 alleviates latency and I/O bottlenecks by providing up to 13x better server performance for web server and e-commerce applications. This performance potential allows data centers to substitute one maxCache 2.0-enabled server solution for thirteen "standard" servers and greatly reduce their CapEx and OpEx.

maxCache 2.0 also allows the potential to replace a SAS HDD configuration with a high-capacity SATA configuration that delivers increased application performance and solution capacity, further enabling data center operators to optimize their hardware investment.

Adaptec by PMC delivers innovative solutions that provide exceptional performance by intelligently routing, optimizing and protecting data as it moves through the I/O path.

With maxCache 2.0 SSD Caching, Adaptec by PMC addresses the business challenges of next-generation data centers and continues to enable the expansion of cloud computing while minimizing environmental and financial costs.



PMC-Sierra, Inc.
1380 Bordeaux Dr.
Sunnyvale, CA 94089 USA
Tel: +1 (408) 239-8000

World Wide Web: www.adaptec.com

Pre-Sales Support: **US and Canada:** 1 (800) 442-7274 or (408) 957-7274 or adaptec-sales@pmc-sierra.com
UK: +44 1276 854 528 or uk_sales@pmc-sierra.com
Australia: +61-2-90116787
Singapore: +65-92351044

Key benefits of maxCache 2.0 include:

Adaptec Series 6Q Storage Controllers with maxCache 2.0 SSD Caching provide the following benefits:

- **Up to 13x faster than HDD-only solutions:** Adaptec maxCache 2.0 SSD Caching software improves the Learned-Path Algorithm that identifies frequently-accessed ("hot") data, and optimizes application performance by copying this data directly into an SSD cache pool for faster retrieval in future requests.
- **Capability to deploy write caching:** Adaptec maxCache 2.0 SSD Caching software offers write caching capability for expanded application workload benefits. Adaptec maxCache 2.0 utilizes the write performance benefits of SSDs to provide additional workload performance advancements.
- **Reduced capital and operating expenses:** Adaptec maxCache SSD Caching reduces capital expenses by increasing IOPs with less hardware – thereby significantly cutting operating expenses related to energy and maintenance.
- **Reduced cost and increased flexibility for SSD selection:** Adaptec 6Q Storage Controllers with maxCache 2.0 SSD Caching allow the use of any standard SSD as a cache, allowing a wide range of vendors and lower cost. These controllers have been qualified with some of the most recent enterprise SSD products to leverage improved performance as well as enterprise features and durability.