



Increase Application Performance Up To 8X With Adaptec maxCache™ SSD Caching Solutions

Executive Summary

Today's data centers and cloud computing environments require increased I/O performance to support large-scale applications such as databases, web servers, e-commerce applications, file servers, and email. This whitepaper focuses on the performance benefits of using Adaptec maxCache SSD Caching Solutions for these key read-intensive applications.

The benefits of maxCache are illustrated by comparing the performance levels between an Adaptec storage controller with maxCache SSD Caching Solution, and an Adaptec storage controller without maxCache SSD Caching Solution. The tests are performed using industry-standard performance tools and/or validation via real world web server application testing.

Performance Highlights

Accelerating application performance, reducing response time, and increasing the number of users supported per server are key requirements for data centers and cloud computing environments. Tests of Adaptec maxCache SSD Caching in web server, e-commerce and mail applications show:

- **Up to 11X performance improvement** in small block random read workloads using Iometer, an industry-standard tool
- **Up to 8X performance improvement** in web server workloads using Iometer
- **Up to 2X performance improvement** in server and workstation application environment
- **Up to 2X the number of users** with improved response time in a real world web server application environment

Introduction

Today's data centers and cloud computing environments require high-density servers to provide end-users with high quality of service. Usually, these server deployments are targeted towards specific applications such as website hosting (for example, web servers) or e-commerce applications (for example, databases and file servers), and accommodate a large number of end-users. Customers of these data centers and cloud computing applications sign service level agreements (SLAs) that obligate the operators to provide pre-determined levels of service, availability, and performance. To meet the service requirements, data center operators deploy servers with high I/O throughput, I/Os per second (IOPS), that are "tuned" to ensure adherence to the SLA. In addition, for Web 2.0 applications, the economies drive to increase the number of hosted users per server to reduce the cost of service (COS) per user.

Growth in the number of end-users translates to an increase in the IOPS required from these servers. Since I/O capacity tends

to run out well before storage capacity, data centers and cloud computing environments must add new servers to keep up with the service requirements, even if the existing servers have the storage capacity for more users. This practice results in low storage capacity utilization in these servers, typically only 30% of total storage available.

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 the 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 products and also sell these "hot" products more often.

Despite having sufficient storage capacity, the servers in these cases are capped in IOPS, and additional servers must be added to keep up with the increased data content. Adding servers increases capital and operating expenses — namely maintenance, power and cooling costs — as well as physical space requirements.

Technology that can provide fast responses for "hot" data while maintaining a large capacity of "less used" 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.

Adaptec maxCache 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.

About Adaptec maxCache SSD Caching Solutions

Solid state drives (SSD) are storage devices capable of high I/O performance. An SSD uses flash components to store data and, unlike a hard disk drive (HDD), has no moving parts and no rotating media.

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 shock 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. SSD's lifespan is also highly dependent on the number of write operations performed to the device.

Adaptec maxCache combines solid state drives 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.

Increase Application Performance 8X With Adaptec maxCache™ SSD Caching Solutions

2

To capitalize on the advantages of SSDs while suppressing their limitations, maxCache SSD Caching introduces a patent-pending Learned-Path Algorithm that identifies frequently-read data, and optimizes “reads” by moving this data directly into an SSD cache for faster retrieval of future requests. By leveraging its unique presence in the data path to create a “read cache pool,” a high-performance hybrid array of HDDs and a single SSD can deliver up to eleven times the IOPS of comparable HDD-only deployments.

Adaptec maxCache represents the next stage in our Data Conditioning Platform strategy, an innovative approach to intelligently routing, optimizing and protecting data as it moves through the I/O path.

Test Methodology

IOMeter was used to compare performance of Adaptec maxCache SSD Caching Solutions using an Adaptec Series 5 and Adaptec Series Q Storage Controller.

Tests were run to measure the benefits of maxCache along two axes. The first set of tests was run to benchmark performance of maxCache across different workloads. A second set of tests was then run to highlight the “scale” benefits of maxCache technology (i.e., increase in performance with increasing number of maxCache enabled SSD).

Workload Performance Test

IOMeter was configured to run with varying queue depths for web server, file server, OLTP (On-Line Transaction Processing), and random read workloads with maxCache SSD caching disabled. The same tests were then run on the same system with maxCache SSD Caching enabled.

Tests consisted of an Adaptec Series 5 controller with Adaptec SSD caching installed on an Intel Xeon platform with a 3.6GHz processor and 2GB system RAM running Microsoft Server 2003 Enterprise Edition 32-bit SP2. To ensure that the base system was set up for best performance, 12 Seagate enterprise-class 15K RPM SAS drives were configured in RAID 5, RAID 0 and RAID 10 configurations. Using SATA drives, and enabling maxCache SSD Caching, will show an even larger performance increase.

Results for Workload Performance Tests

The test results highlighted performance increases of **up to 11X** with Adaptec SSD cache enabled. Figure 1 illustrates the following benefits of Adaptec SSD caching:

- **Up to 11X performance improvement** for small block random read workloads
- **Up to 8X performance improvement** for web server workloads
- **Up to 8X performance improvement** for file server workloads
- **Up to 5X performance improvement** for OLTP workloads

Web server and e-commerce applications have high throughput requirements and are read-intensive. Random read, web server and file server workloads best represent these data-rich applications with predominantly read transactions and interspersed writes.

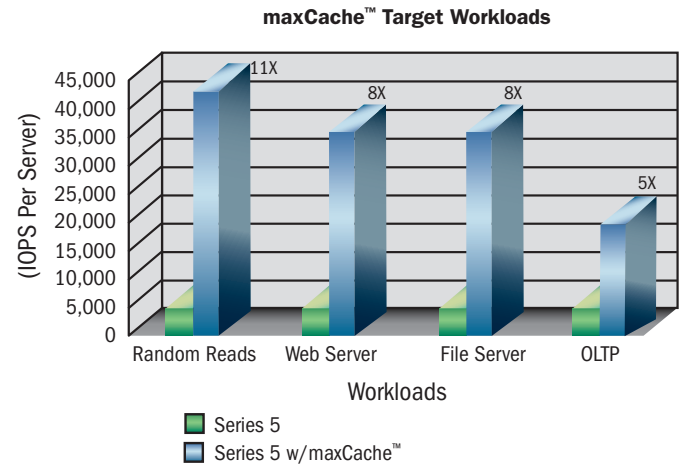


Figure 1

Scalability Test

IOMeter was then set up to run with file server and web server workloads (i.e., workloads that represent web hosting and e-commerce applications). The test was run with maxCache SSD Caching disabled then repeated on the same system with maxCache SSD Caching enabled with one, and then two SSDs used as a cache pool. Data sets were also scaled according to the size of the total SSD cache to ensure a fair comparison of performance.

Test setup consisted of an Adaptec maxCache 5805Q controller installed on an Intel Xeon processor 3.6GHz system containing 2GB system RAM running Microsoft Server 2003 Enterprise Edition 32-bit SP2. To ensure that the base system was set up for best performance, 12 Seagate enterprise class 15K RPM SAS drives were configured in a RAID 0 configuration.

Performance Results for Scalability Tests

Test results show that I/O throughput of the server system scales as more SSDs are added to the server as cache. As shown in Figure 2 and in Figure 3 for both web server and file server workloads, I/O performance scales beyond 7X with the addition of SSDs used as cache.




Figure 2: Using Adaptec 5805Q, with 100GB SSDs, 12 SAS HDDs, at 512K block size

Increase Application Performance 8X With Adaptec maxCache™ SSD Caching Solutions




Figure 3: Using Adaptec 5805Q, with 100GB SSDs, 12 SAS HDDs, at 512K block size

Web Server Application Test

The true measure of a product’s benefit is observed during real-world usage. Further proof of maxCache’s performance superiority is found in the real-world application testing of a web server setup conducted by AppLabs, an independent performance testing company. The goal of the real world testing was to determine the additional number of users hosted by adding maxCache SSD Caching to the web server.

Test setup consisted of an Adaptec Series 5 controller with Adaptec SSD Caching installed on an Intel 3.6GHz Xeon platform containing 2GB system RAM running Red Hat 5 installed on the RAID volume. To emulate a typical server used in data center and cloud computing environments, the base system was set up with three Western Digital SATA drives in a RAID 5 configuration. As shown in Figure 4, the setup consists of multiple users generating traffic to a web server configured with maxCache SSD Caching.




Figure 4

The Web server was set up to host a total of 8,300 websites with multiple pages constituting over 40GB of data. Traffic was generated by emulating clicks on the websites from multiple users.

Aberdeen Group research shows that business performance begins to suffer after 5.1 seconds delay in response times of web applications. An additional 1 second of delay adversely impacts customer satisfaction by up to 16%. These benchmarks were incorporated into AppLabs’ web server application test.

The test was first run with maxCache SSD Caching disabled. Throughput, transaction rate and response times were measured while increasing the number of users until the web server response time exceeded 5.1 seconds. The test was then repeated with maxCache SSD Caching enabled to capture the increase in number of users, reduction in response times, and the increased throughput supported by the web server. After reaching peak user load in each of the tests, the system was then run for 45 minutes.

This real world testing proved maxCache’s superior performance and showed that maxCache was able to support twice the number of users on the same web server system while improving response times by up to 5X, as shown in Figure 5.

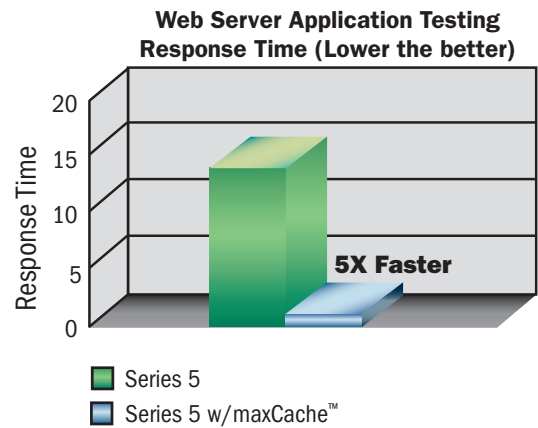


Figure 5

As shown in Figure 6 and Figure 7, the transaction completion rate and throughput of the web server increased by 4X when maxCache SSD Caching was enabled. The tests also showed that the learning algorithm identified the most frequently-read data and cached it for future use, further increasing the throughput of the web server in the maxCache-enabled mode.




Figure 6

Increase Application Performance 8X With Adaptec maxCache™ Caching Solutions




Figure 7

Enhanced IOPs Performance with Random Reads

Figure 8 shows an impressive 12X improvement in IOPs performance using the Adaptec maxCache 5805Q controller with four Samsung 100GB SSDs in the read cache.




Figure 8: Using Adaptec 5805Q, with 100GB SSDs, 12 SATA HDDs, at 512K block size

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-efficient products to enhance the server's I/O capability.

As highlighted in the tests above, maxCache alleviates the I/O bottleneck by providing up to 8X better server performance for web server and e-commerce applications.

Therefore, instead of adding eight additional servers to keep up with the increasing throughput demands, data center operators can add one Adaptec maxCache SSD Cache Performance Solution for a fraction of the cost.

Key benefits of maxCache include:

- **Accelerated application performance** by up to 8X more than HDD-only arrays.
- **Reduced capital and operating expenses by up to 50%** by increasing the number of users per server and reducing the number of servers needed to provide end-users with the pre-determined service quality levels.
- **Improved server & storage utilization** via a patent-pending Learned-Path Algorithm that identifies frequently-read “hot” data, and optimizes “reads” by copying this data directly into an SSD cache for faster retrieval of future requests.
- **Seamless compatibility** with today's off-the-shelf servers, storage devices and operating systems. Adaptec maxCache drops right into all industry-standard platforms, and provide application agnostic caching that does not require changes to storage architectures, application software or operating systems.
- **Flexible SSD cache size** can be easily expanded by adding up to four maxCache SSDs per Storage Controller. Adaptec maxCache devices can also be removed or replaced anytime without impacting data integrity since all user data is permanently stored on the hard disk drive or arrays.

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

With the introduction of the maxCache SSD Cache Performance Solution, Adaptec by PMC helps address the business needs of next-generation data centers and cloud applications, while reducing capital equipment and operating costs.



PMC-Sierra, Inc.

Mission Towers
3975 Freedom Circle
Santa Clara, CA 95054 USA
Tel: +1 (408) 239-8000
Fax: +1 (408) 492 9192

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: +61-2-95031555
Singapore: +65-92351044

© Copyright PMC-Sierra, Inc. 2010. All rights reserved. PMC, PMC-SIERRA and Adaptec are registered trademarks of PMC-Sierra, Inc. “Adaptec by PMC” is a trademark of PMC-Sierra, Inc. Other product and company names mentioned herein may be trademarks of their respective owners. For a complete list of PMC-Sierra trademarks, see www.pmc-sierra.com/legal.

Part Number: 667140-011-4.0 Printed in U.S.A. 10/10