Microsemi

Power Matters.™

maxCache 3.0 SSD Read and Write Caching Solutions

Microsemi® Adaptec® Series 8Q/8QZ (12 Gbps) and Series 7Q White Paper

March 2016

Microsemi
Power Matters.™

SSD Read and Write 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 Microsemi's Adaptec maxCache 3.0 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.

In terms of I/O performance and average latency, SSDs can deliver up to a 100x performance gain over Hard Disk Drives (HDDs) in random-access read and write operations. maxCache leverages those benefits by using SSDs to cache copies of frequently-accessed (aka "hot") data for both read and write workloads.

The performance and financial benefits of maxCache 3.0 are illustrated in this paper by comparing the performance levels in a 100% Random-Read IOmeter workload scenario with a maxCache 3.0-enabled Microsemi Adaptec RAID adapter vs. the same adapter without maxCache 3.0 enabled, as well as an OLTP scenario with a maxCache 3.0-enabled Microsemi Adaptec RAID adapter vs. an adapter without maxCache 3.0 enabled. 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).

maxCache

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.

Benchmark tests designed to illustrate the theoretical maximum benefits of maxCache 3.0 examined a scenario where all data was delivered by the maxCache SSD cache (a 100% hit rate) and compared the results to those of an HDD-only RAID array of eight 10k RPM SAS HDDs. The comparison showed:

- **Up to 44x performance improvement** in read-intensive I/O operations per second (IOPS) in a RAID 0 configuration
- **Up to 31x performance improvement** in read-intensive IOPS in a RAID 5 configuration
- **Up to 44x reduction in latency** in read-intensive applications in a RAID 0 configuration
- **Up to a 31x reduction in latency** in read-intensive applications in a RAID 5 configuration

In a typical real-world environment where the maxCache scenario incorporates data delivered by both the SSDs and HDDs, the maxCache solution still delivered substantial benefits over the HDD-only array:

- **4x performance improvement** in mixed workload IOPS in a RAID 0 configuration
- **4x performance improvement** in mixed workload IOPS in a RAID 5 configuration
- **4x latency reduction in mixed workload IOPS** in a RAID 0 configuration
- **4x latency reduction in mixed workload IOPS** in a RAID 5 configuration
- **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 per user. 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.When more Website visitors hurt your business: Are you ready for peak traffic?,” Equation Research 2010

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.

Microsemi's Adaptec maxCache 3.0 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.

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

About maxCache SSD Caching Solutions

Microsemi Adaptec maxCache combines SSDs used as cache with Microsemi'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 3.0 SSD caching can provide significant performance gains compared to HDD-only deployments (Figure 1).

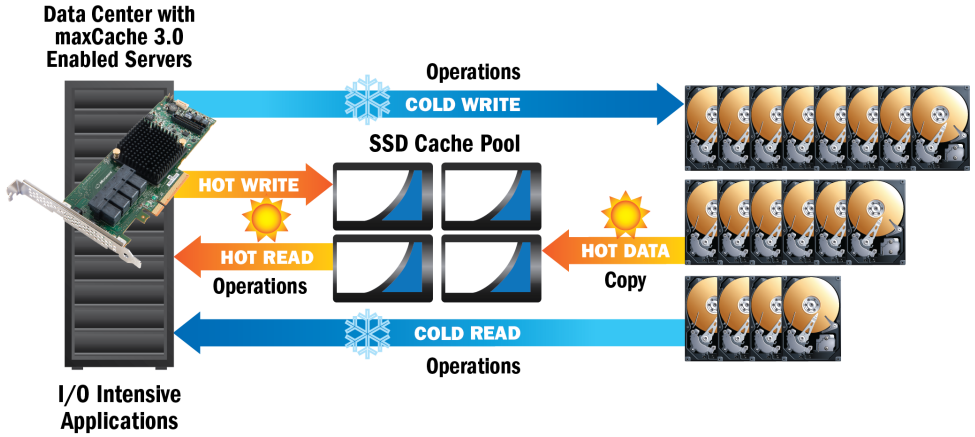

Figure 1 • **Microsemi Adaptec maxCache 3.0 Deployment**

Advancements in maxCache 3.0's Learned-Path Algorithm have improved the read caching function to support non-redundant cache pool (RAID 0) scenarios. In this case, an SSD failure will not impact data availability as long as the HDDs are configured in a redundant RAID since all data is still securely stored on the HDD RAID.

maxCache 3.0 has also evolved to include write-back caching. By caching writes to a redundant SSD cache pool (either RAID 1E or RAID 5), maxCache 3.0 leverages the performance and latency capabilities of SSD technology for both read and write workloads.

By expanding the use of SSD caching, maxCache 3.0 is suitable for wide scale data center deployments, offering increased financial benefit while streamlining hardware design and implementation.

Also new to maxCache 3.0 is the inclusion of Optimized Disk Utilization (ODU). As SSD storage sizes continue to grow, monopolizing their entire capacity as a cache-pool is becoming less attractive. ODU allows the SSD to be partitioned into both a cache pool and logical device (Figure 2). The logical partition, unlike the cache pool, is exposed to the operating system, and can be used as an OS boot drive, or to store other data that requires fast, low-latency access.

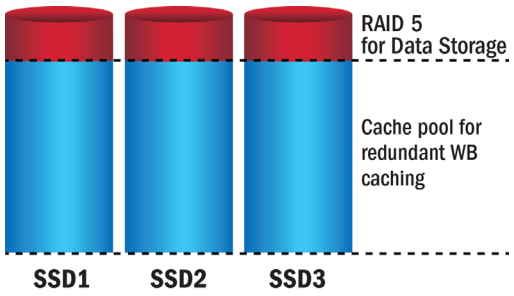

Figure 2 • **Microsemi Adaptec maxCache 3.0 Deployment**

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.

maxCache 3.0 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 3.0 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 an administrator to have intimate knowledge of that specific application to manually tune it to route certain data to the SSDs.

maxCache eliminates this manual application-specific tuning by automatically and transparently analyzing and routing the read and write data.

Additionally, maxCache 3.0 offers flexibility that PCIe-based flash cards do not. With maxCache 3.0, the end-user can install an SSD with performance metrics and write endurance best suited for the application. If a maxCache SSD fails, it can be easily replaced. Conversely, if a portion of the flash fails in a PCIe-based scenario, the entire flash card must be replaced.

Recommended SSDs for maxCache 3.0

Due to the sheer volume of data being written to the SSDs in write-caching scenarios, maxCache 3.0 requires enterprise class SSDs. Compared to the client class SSDs that can be found in personal systems such as ultrabooks, laptops and desktop computers, enterprise class SSDs offer superior write endurance (the number of write cycles a block of flash memory can accept before it becomes unusable), support heavier write activity (assume 24 hours per day every day for a data center vs. 8 hours on weekdays for an employee's personal computer), and function in more extreme environmental conditions.

For read-caching only scenarios, the amount of writes are limited. While enterprise class SSDs are still recommended, more cost efficient SSDs can also be used, such as those with less over provisioning and lower write endurance.

Test Methodology

IOmeter was used to measure the performance benefits of maxCache 3.0 SSD Caching Solutions by comparing the base combination of a Microsemi Adaptec Series 7Q RAID Controller and HDDs to the same combination with SSDs and maxCache added.

One set of tests was run in a 100% read-only workload scenario, and other set was run in a mixed read and write workload scenario. Both sets of tests were run in both RAID 0 and RAID 5 configurations.

Performance

Read Caching Performance — RAID 0

Read-intensive applications such as web servers, file servers, and e-commerce applications radically benefit from maxCache 3.0, yielding significant performance gains in IOPS.

In RAID 0, maxCache delivered a 44x increase in IOPS compared to HDD-only arrays.

- RAID 0 performance comparison under 100% Random Read IOmeter workload (Figure 3).
- HDD-only configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500 GB capacity each, in RAID 0.
- maxCache 3.0 configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500 GB capacity each; six 6 Gbps SATA SSDs, 50 GB each, in RAID 0.

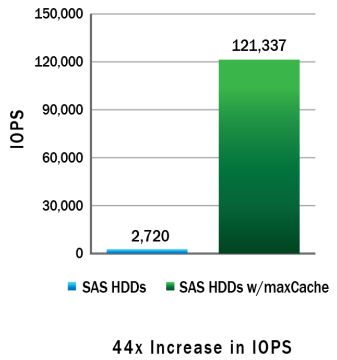

Figure 3 • **44x Increase in IOPS with maxCache 3.0**

Read Caching Performance — RAID 5

In RAID 5, maxCache yielded a 31x increase in IOPS compared to HDD-only arrays (Figure 4).

- RAID 5 performance comparison under 100% Random Read IOmeter workload.
- HDD-only configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500GB capacity each, in RAID 5.
- maxCache 3.0 configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500 GB capacity each; six 6 Gbps SATA SSDs, 50 GB each, in RAID 5.

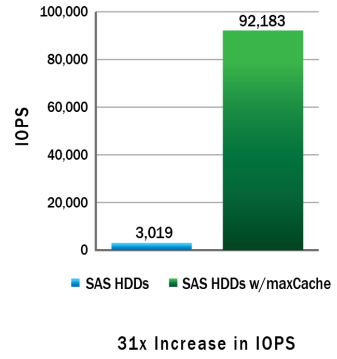

Figure 4 • **31x Increase in IOPS with maxCache 3.0**

Write Caching Performance

Improved write caching support extends maxCache 3.0 benefits to I/O-intensive applications with mixed workloads, including OLTP, Microsoft Exchange Server, and HPC environments. Enabling maxCache 3.0 in these scenarios delivers substantial performance benefits while also allowing for a server design with much higher capacity potential.

These performance gains were achieved by eliminating a SAS HDD server configuration in favor of a maxCache 3.0 configuration with high-capacity SATA HDDs.

The configuration change expanded the available solution capacity from 2.4 TB to 18 TB – a 7x increase that significantly reduces a data center's cost per GB and cost per I/O metrics.

Write Caching Performance — RAID 0

In RAID 0, maxCache yielded a 4x increase in IOPS compared to HDD-only arrays (Figure 5).

- RAID 0 performance comparison under 100% Random Read IOmeter workload.
- HDD-only configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500GB capacity each, in RAID 0.
- maxCache 3.0 configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500 GB capacity each; six 6 Gbps SATA SSDs, 50 GB each, in RAID 0.

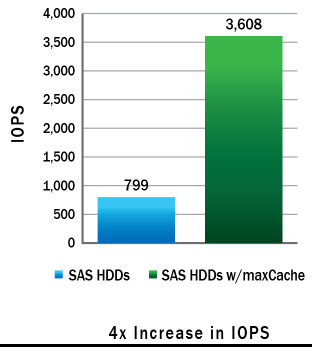

Figure 5 • **4x Increase in IOPS with maxCache 3.0**

Write Caching Performance — RAID 5

In RAID 5, maxCache yielded a 4x increase in IOPS compared to HDD-only arrays (Figure 6).

- RAID 5 performance comparison under 100% Random Read IOmeter workload.
- HDD-only configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500 GB capacity each, in RAID 5.
- maxCache 3.0 configuration: Ten 7200 RPM 6 Gbps SAS HDDs, 500GB capacity each; six 6 Gbps SATA SSDs, 50 GB each, in RAID 5.

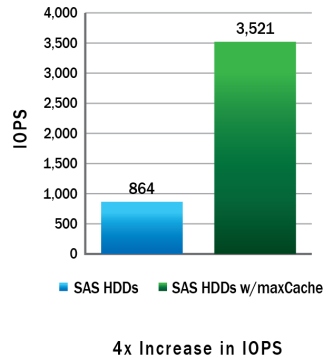

Figure 6 • **4x Increase in IOPS with maxCache 3.0**

Impact and Savings

Real-World Cost-Savings

By accelerating IOPS and reducing latency, maxCache 3.0 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 3). 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.

Real-World Impact of maxCache on CapEx and OpEx

The above test data represent the theoretical maximum performance benefits in scenarios where 100% of the data is being read from or written to the SSD cache. In the real world, however, only a fraction of data is "hot" and can be read from or written to the SSD cache. Therefore, most applications will see benefits that are less than the maximum capabilities discussed above.

Environments

Read-only Environments

To calculate the cost-saving benefits of a typical real-world maxCache 3.0 environment, we will assume that 90% of the data is "hot" (delivered by the SSDs), and 10% is "cold" (delivered by the HDDs).

Under these conditions, the maximum performance of the application is determined by the amount of time needed for the slower HDDs to process their 10% of the data.

We assume that SSDs are 9x faster than HDDs on random-access reads, so the maxCache 3.0 solution would yield a 9x performance benefit on hot data reads, compared to a HDD-only array.

Mixed Workload Environments

In mixed workload RAID 5 and RAID 0 scenarios, the real-world performance improvements presented earlier show that maxCache 3.0 delivers 4x greater IOPS and 4x less latency than HDD-only arrays.

Therefore, maxCache 3.0 allows data center and cloud computing environments to host the same number of users and perform the same number of transactions per second on a quarter of its current HDD-only server infrastructure.

These performance gains significantly reduce 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 7). 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.

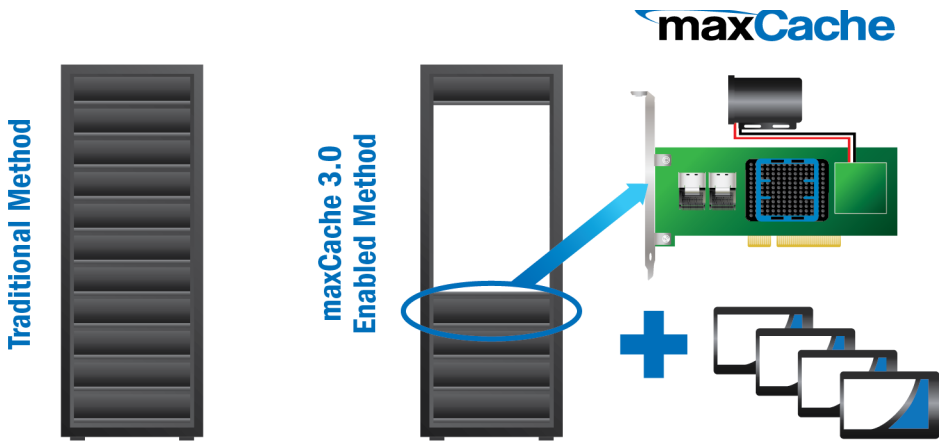

Figure 7 • **Qualified Impact on Utilization of Hardware Investment**

Notes:

* 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.

Given the performance gains demonstrated above, it may seem preferable to build an SSD-only RAID array instead of a mixture of SSDs and HDDs. Keep in mind, however, that standard SSD storage capacities are still relatively small compared to those of HDDs, making their cost per GB much higher. Adding too many SSDs would significantly increase CapEx.

maxCache 3.0 delivers an optimized balance of SSD performance and HDD capacity to solve the unique challenges faced by data centers and cloud computing operators. With maxCache 3.0 SSD Caching, a data center can still benefit from the larger capacities of rotating media (HDDs), while gaining the improved I/O performance benefits of SSDs.

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 3.0 alleviates latency and I/O bottlenecks by providing up to 44x better server performance for web server and e-commerce applications. This performance potential allows data centers to substitute one maxCache 3.0-enabled server solution for thirteen "standard" servers and greatly reduce their CapEx and OpEx.

maxCache 3.0 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.

Microsemi delivers innovative solutions with Series 7 and 8 products that provide exceptional performance by intelligently routing, optimizing and protecting data as it moves through the I/O path.

With maxCache 3.0 SSD Caching, Microsemi addresses the business challenges of next-generation data centers and continues to enable the expansion of cloud computing while minimizing environmental and financial costs.

Key Benefits of maxCache 3.0

Microsemi Adaptec Series 7Q Storage Controllers with maxCache 3.0 SSD Caching provide the following benefits:

- **Up to 44x faster than HDD-only solutions**: maxCache 3.0 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**: maxCache 3.0 SSD Caching software offers write caching capability for expanded application workload benefits. maxCache 3.0 utilizes the write performance benefits of SSDs to provide additional workload performance advancements.
- **Reduced capital and operating expenses**: 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**: Microsemi Adaptec 7Q Storage Controllers with maxCache 3.0 SSD Caching allows the use of any enterprise class 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.

Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996

E-mail: sales.support@microsemi.com

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for communications, defense & security, aerospace and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, Calif., and has approximately 4,800 employees globally. Learn more at **www.microsemi.com**.

Document No.: ESC-2160398 , Issue 1