Microsemi Adaptec® Hybrid RAID Solutions

Combining SSDs and HDDs for Maximum Performance and Reliability

White Paper

Microsemi
Power Matters.™

Introduction

As user demand for speed grows and applications become more and more robust, data centers face a constant battle to increase their infrastructure performance accordingly. Some components offer a relatively straightforward upgrade path: CPUs continue to get more powerful, networking speeds increase, disk drive capacities grow, and so on. But these advances mean little if the rate at which I/O operations can be performed doesn't keep pace.

For example, an online transaction-processing system is often limited by the rate at which database updates can occur. The performance of a search algorithm may depend on how quickly various data structures can be read from disk. The number of users that can successfully access a website depends on how quickly the web pages can be served.

Growth in the number of end-users compounds the problem by causing longer response times and increasing latency. A study by Equation Research illustrates the impact such issues can have on a company's website[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 with a less positive perception of the company.

To boost system performance, data centers will sometimes add very fast hard disk drives (HDDs)—for example, 15,000 RPM Serial Attached SCSI (SAS) drives—and "short stroke" them, which means that applications use only the outermost tracks on the disks where the performance is best. Or, they may add new servers, even if the existing servers have the storage capacity for more users.

Either method, however, results in low storage-capacity utilization. On top of that, adding more servers leads to an increase in capital and operating expenses—namely maintenance, power, and cooling costs—as well as physical space requirements.

The proliferation of solid-state drives (SSDs) alleviates many of these problems, and changes the way data centers think about their storage architecture.

SSDs: Performance in a Flash

SSDs offer several advantages over HDDs, including higher read bandwidth, higher IOPs, better mechanical reliability, and higher resistance to shock and vibrations. However, the same features (such as flash) that provide these advantages also come with some inherent limitations compared to HDDs, such as limited capacity and lower streaming-write bandwidth.

Another limitation of SSDs is that they cannot directly overwrite data like an HDD can. Instead, the entire destination data block must first be erased before new data can be written—even if only a few bytes are changed. Not only does this process negatively impact write performance, it can cause uneven wear on the surface of the flash medium and lead to premature failure of the drive.

To combat uneven wear, many SSDs employ wear leveling to distribute erasures and rewrites evenly across the medium surface. Wear leveling extends the life of the drive, but also further degrades write performance over time.

---

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

SSDs are also much more expensive than HDDs, especially when measured using the traditional "cost per GB of capacity" metric. In most environments it is simply impractical to replace every HDD with an SSD, regardless of the potential benefits.

However, a hybrid approach that combines HDDs and SSDs can be extremely cost efficient, especially when measured in terms of the "cost per I/O operation" or the more green-oriented "power consumption per I/O operation".

What is Hybrid RAID?

Hybrid RAID is a redundant storage solution that combines high-capacity, low-cost SATA or higher-performance SAS HDDs with low-latency, high-IOPs SSDs, and an SSD-aware RAID adapter card, as shown in the following illustration.

Figure 1 Hybrid RAID Deployment

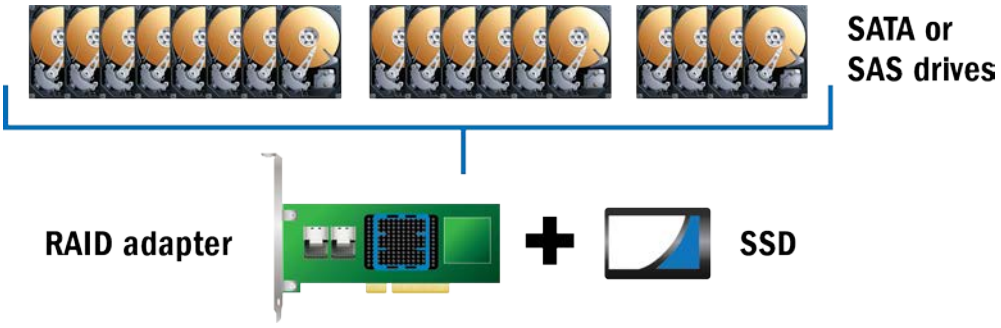

In hybrid RAID, read operations are done from the faster SSD, and write operations happen on both SSD and HDD for redundancy purposes.

Hybrid RAID is ideal for applications with low levels of data, such as an Internet gateway, file server, or virtual machine.

What Are the Benefits of Hybrid RAID?

Hybrid RAID arrays offer tremendous performance gains over standard HDD RAID arrays at a much lower cost than SSD-only RAID arrays. Compared to HDD-only RAID arrays, hybrid arrays accelerate IOPs and reduce latency, allowing data center and cloud-computing environments to host more users and perform more transactions per second on each server, which reduces the number of servers required to support any given workload.

Not only does the reduction in servers lead to reduced server footprint requirements within the data center, there is the additional financial benefit of reducing both the capital expenses required to purchase additional servers, as well as the operating expenses associated with power, cooling, and maintenance.

Hybrid RAID Examples

Incorporating hybrid RAID into a storage solution opens up a wide range of possibilities for maximizing capacity, boot times, and overall performance, as shown in the following examples.

Scenario 1: Large-Capacity Servers (Ideal for Series 8)

In this scenario, the customer wants to make the maximum use of space in a large-capacity, 8-bay server, while still gaining a fast-booting server. The customer is conservative and wants to keep the OS away from data (OS in mirror and data in RAID 5 for maximum capacity).

Figure 2 Large-Capacity Server, Traditional HDD-Only Method

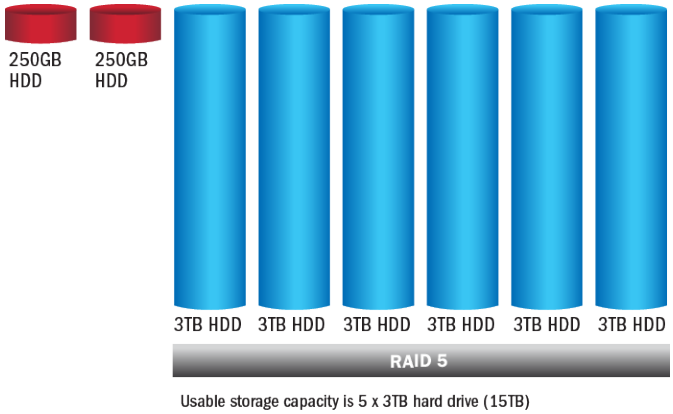

The problem with an HDD-only method, as shown in the previous illustration, is that it wastes two bays in the server to the OS boot drives, and therefore limits capacity to six drives in RAID 5. Additionally, the server does not boot particularly quickly as it is booting from standard HDD in mirror, which is not a particularly fast form of RAID.

Figure 3 Large-Capacity Server, Hybrid RAID Solution

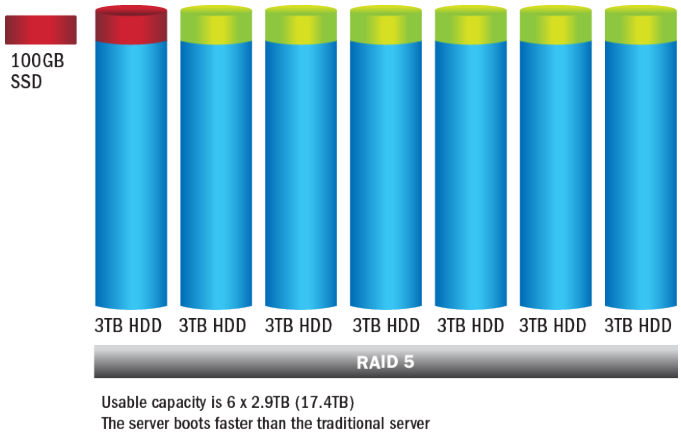

The hybrid solution, as shown in the previous illustration, offers several benefits:

1. Total server capacity is six times the usable capacity of each drive in a RAID 5.
2. The server will boot quickly, as it is reading from an SSD during the boot process.
3. The unused disk space on each disk can be utilized in another array if required.
4. In the above scenario, a 500 GB RAID 5 disk could be created across the six unused disk segments.

Scenario 2: Workstation (Ideal for Series 8 and Series 8E)

In this scenario, the workstation user wants a fast-booting, fast application-loading system, and doesn't want the hassle of having to rebuild the workstation if a drive fails.

Figure 4 Workstation Hybrid RAID Solution

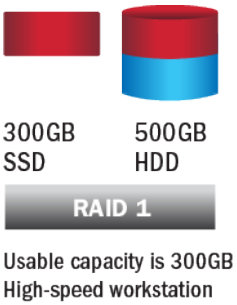

The hybrid solution, as shown in the previous illustration, provides the following advantages:

- All reads are directed from the SSD, so read speed of the workstation is dramatically improved over that of a standard single HDD.
- Write speed is the same as a standard workstation, but due to the cache on the adapter card, it will be faster than a standard HDD connected to an on-board disk adapter.
- Data is not lost if either drive fails. When the failed drive is replaced, the array will rebuild onto the replacement drive, putting the system back the way it was before the drive failure.

Note that it is not recommended to use this unused 200 GB capacity on the 500 GB HDD due to the fact that this data is not redundant, and failure of the HDD would result in data loss.

Scenario 3: Small Business Server (Ideal for Series 8E)

Many small business servers require at least some high-performance storage component to handle accounting software, industry-specific small databases, and even mail servers. Along with that, performance is always the requirement for capacity. Even small businesses can easily generate several TB of data in the form of documents, photos, video, etc.

Figure 5 Small Business Server, Hybrid RAID Solution

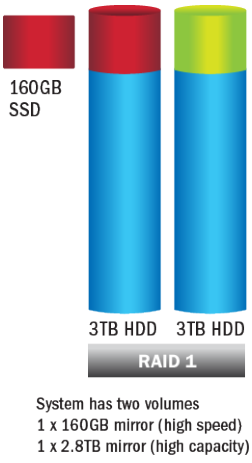

This hybrid solution provides the following advantages:

- The 160 GB disk is big enough to create a boot volume for the OS while still leaving enough space for a 100 GB volume for database function.
- Because the 160 GB hybrid mirror reads directly from the SSD, both OS boot time and database function will be greatly improved over a standard mirror of two HDDs.

Scenario 4: Small Business Server (Ideal for Series 8E)

The previous configuration requires at least a 4-port RAID card, but the fourth port is unused. Extending the server to utilize all four ports is simply a matter of purchasing a second 160 GB SSD to mirror to the unused space from the previous example.

Figure 6 Small Business Server, Hybrid RAID Solution

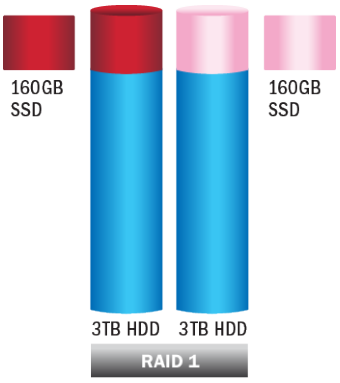

It might seem more sensible to mirror the two SSDs, and then mirror the two HDDs in what would seem a more conventional server. However, there is good reason to do the above:

- The write speed of lower-spec SSDs is not much faster than that of HDDs, so there is little to be gained by mirroring the two SSDs.
- Since the read speed of the SSD is so quick, putting them in a mirror will not result in much improvement over a single SSD.
- Since the focus is on read speed, having two Hybrid RAID arrays now gives 320 GB of SSD read speed to the server, while still protecting the data on each SSD in the form of a mirror.
- The HDD capacity of the server is not sacrificed from the previous example.

Scenario 5: Small Business Server, High Speed, Small Capacity (Ideal for Series 8E)

If performance is more important than capacity, then the user can install two 600 GB SSDs and mirror them to the HDD, thus gaining 1.2 TB of SSD capacity, as shown in the following illustration.

Figure 7 Small Business Server, High Speed, Small Capacity Hybrid RAID Solution

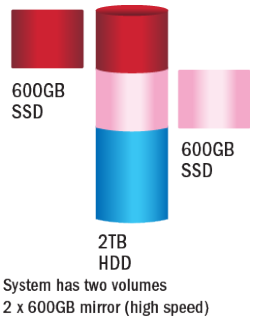

Note that it would be possible in this configuration to add a third SSD (using the fourth port on the RAID card), to give three hybrid mirrors on a total of three SSDs and one HDD.

Hybrid RAID Solutions

Building a Hybrid Solution

From a hardware perspective, building a hybrid RAID solution is fairly straightforward as any capacity of SSDs and HDDs can be used (though an equal number of SSDs and HDDs must be used). If the RAID array is built from different-sized drives, the drive segment size is the size of the smaller disk drive. For example, a RAID 1 created with one 128 GB SSD and one 2 TB HDD will make a 128 GB logical device. A RAID 10 created with two 128 GB SSDs and two 2 TB HDDs will make a 256 GB logical device. The remaining HDD capacity can be used for storage.

From an application perspective, though, it is not quite as simple since most software was not created to be aware of the possibility of having two types of storage with such different characteristics.

In order to take full advantage of a hybrid RAID setup, an SSD-aware RAID adapter with intelligent storage processing must be deployed.

Microsemi Adaptec Hybrid RAID Solutions

Fortunately, reaping the rewards of hybrid RAID technology can be as simple as installing updated firmware onto any Series 8, Series 8E, Series 7, Series 6, or Series Q SAS/SATA RAID adapter from Microsemi. The adapter takes care of the rest by automatically creating a hybrid RAID array whenever one or more SSDs are combined with the same number of HDDs in a RAID 1 or RAID 10 array. The hybrid RAID is completely transparent to the operating system and all running applications.

Additionally, Microsemi Adaptec RAID adapters deliver maximum hybrid array performance by writing to both HDD and SSD, and reading from SSD 100% of the time, as shown in the following illustration.

Figure 8 Microsemi Adaptec Hybrid RAID Solution

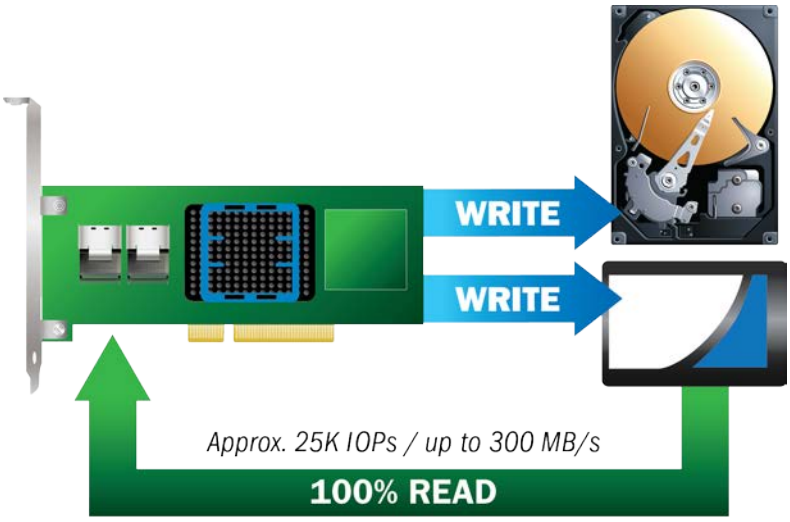

By contrast, other adapters write and read from HDD and SSD, as shown in the following illustration. Since server I/O of an SSD can be up to 100 times faster than HDD, competitive hybrid RAID solutions typically deliver lower I/O performance.

Figure 9 Competitive Hybrid RAID Solution

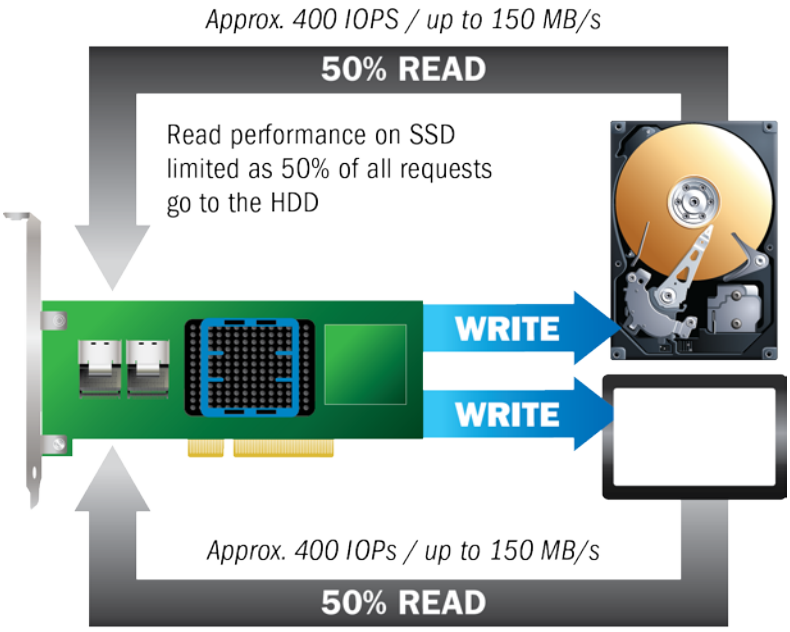

Microsemi Adaptec hybrid RAID arrays deliver a higher number of read operations per second than standard HDD arrays with no degradation of write I/O performance, as shown in the following illustration.

Figure 10 IOMeter Raw Performance

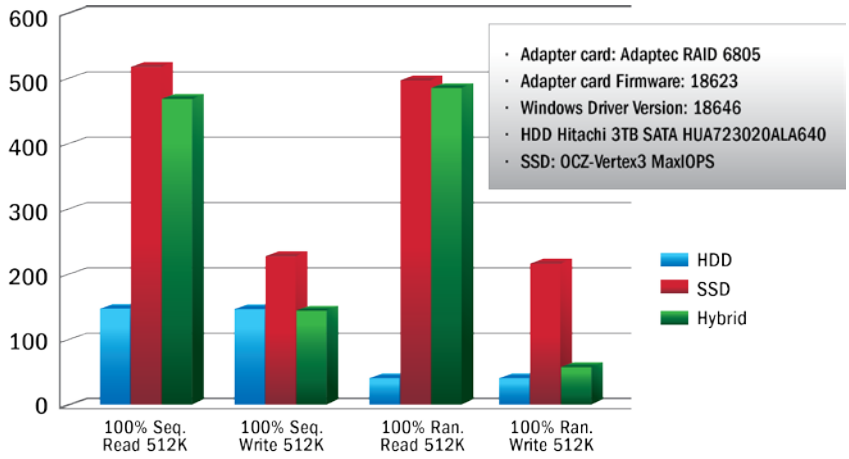

Some applications even see a slight performance benefit from Microsemi Adaptec hybrid RAID and are faster on a Hybrid RAID 1 than on a single SSD, as shown in the following illustration.

Figure 11 AS SSD Application Performance

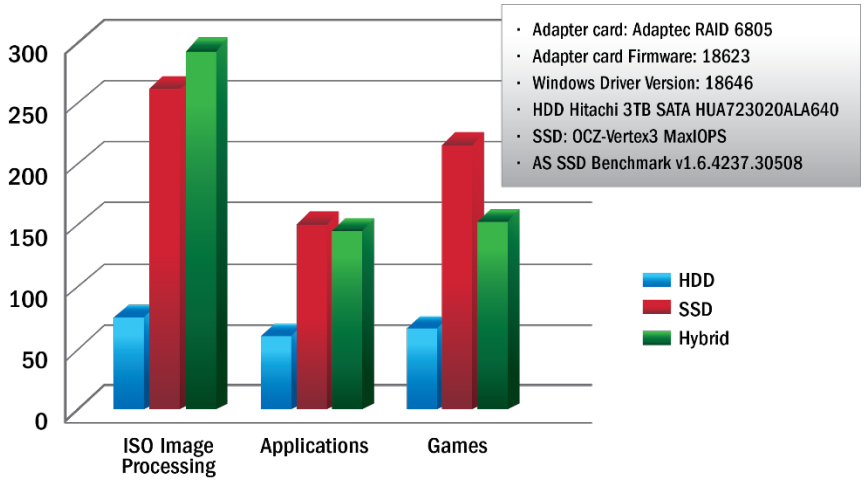

As such, Microsemi Adaptec hybrid RAID offers better performance and lower cost per GB, and lower cost and power consumption per I/O than competing hybrid RAID solutions.

Conclusion

Data centers face never-ending pressure to accommodate more users and more traffic. In the face of tightening budgets, however, the days of simply adding more servers are coming to an end.

Hybrid RAID arrays of SSDs and HDDs offer significantly better performance than standard HDD RAID arrays at a lower cost than SSD-only arrays.

Microsemi's Adaptec Series 8, Series 8E, Series 7, Series 6, and Series Q SAS/SATA RAID adapters with the latest firmware updates can automatically create a seamless Hybrid RAID array whenever one or more SSDs are combined with the same number of HDDs in a RAID 1 or RAID 10 array. Unlike

competing RAID adapter cards that read from the HDD 50% of the time, Microsemi Adaptec cards deliver maximum system performance by reading only from the SSD.

Compared to HDD-only arrays, Microsemi Adaptec hybrid arrays deliver a higher number of read operations per second with no degradation of write I/O performance, and complete transparency to the operating system and all running applications. Compared to SDD-only arrays, Microsemi Adaptec hybrid arrays offer the lowest cost per GB and lowest cost and power consumption per I/O.

Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA
Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Fax: +1 (949) 215-4996
Email: sales.support@microsemi.com
www.microsemi.com

About Microsemi

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center 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; enterprise storage and communication solutions; 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, California, and has approximately 4,800 employees globally. Learn more at www.microsemi.com.

ESC-2160436