



Choosing the Right RAID

Which RAID Level is Right for You?

White Paper

April 2016

Choosing the Right RAID

Introduction

For any organization, whether it be small business or a data center, lost data means lost business. There are two common practices for protecting that data: backups (protecting your data against total system failure, viruses, corruption, etc.), and RAID (protecting your data against drive failure). Both are necessary to ensure your data is secure.

This white paper discusses the various types of RAID configurations available, their uses, and how they should be implemented into data servers.

Note: RAID is not a substitute for regularly-scheduled backups. All organizations and users should always have a solid backup strategy in place.

What is RAID?

RAID (Redundant Array of Inexpensive Disks) is a data storage structure that allows a system administrator/designer/builder/user to combine two or more physical storage devices (HDDs, SSDs, or both) into a logical unit (an array) that is seen by the attached system as a single drive.

There are three basic RAID elements:

1. Striping (RAID 0) writes some data to one drive and some data to another, minimizing read and write access times and improving I/O performance.
2. Mirroring (RAID 1) replicates data on two drives, preventing loss of data in the event of a drive failure.
3. Parity (RAID 5 & 6) provides fault tolerance by examining the data on two drives and storing the results on a third.
4. When a failed drive is replaced, the lost data is rebuilt from the remaining drives.

It is possible to configure these RAID levels into combination levels — called RAID 10, 50 and 60.

The RAID controller handles the combining of drives into these different configurations to maximize performance, capacity, redundancy (safety) and cost to suit the user needs.

Hardware RAID vs. software RAID

RAID can be hardware-based or software-based.

Hardware RAID resides on a PCIe controller card, or on a motherboard-integrated RAID-on-Chip (ROC). The controller handles all RAID functions in its own hardware processor and memory. The server CPU is not loaded with storage workload so it can concentrate on handling the software requirements of the server operating system and applications.

Pros:

- Better performance than software RAID.
- Controller cards can be easily swapped out for replacement and upgrades.

Cons:

- More expensive than software RAID.

Software RAID runs entirely on the CPU of the host computer system.

Pros:

- Lower cost due to lack of RAID-dedicated hardware.

Cons:

- Lower RAID performance as CPU also powers the OS and applications.

How does RAID work?

In software RAID, the RAID implementation is an application running on the host. This type of RAID uses drives attached to the computer system via a built-in I/O interface or a processor-less host bus adapter (HBA). The RAID becomes active as soon as the OS has loaded the RAID driver software.

In hardware RAID, a RAID controller has a processor, memory and multiple drive connectors that allow drives to be attached either directly to the controller, or placed in hot-swap backplanes.

In both cases, the RAID system combines the individual drives into one logical disk. The OS treats the drive like any other drive in the computer — it does not know the difference between a single drive connected to a motherboard or a RAID array being presented by the RAID controller.

Given its performance benefits and flexibility, hardware RAID is better suited for the typical modern server system.

RAID-compatible HDDs and SSDs

Storage manufacturers offer many models of drives. Some are designated as “desktop” or “consumer” drives, and others as “RAID” or “enterprise” drives. There is a big difference: a consumer drive is not designed for the demands of being connected into a group of drives and is not suitable for RAID. RAID or enterprise drives, on the other hand, are designed to communicate with the RAID controller and act in unison with other drives to form a stable RAID array to run your server.

From a RAID perspective, HDDs and SSDs only differ in their performance and capacity capabilities. To the RAID controller they are all drives, but it is important to take note of the performance characteristics of the RAID controller to ensure it is capable of fully accommodating the performance capabilities of the SSD. Most modern RAID controllers are fast enough to allow SSDs to run at their full potential, but a slow RAID controller could bottleneck data and negatively impact system performance.

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 (Figure 1). 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 arrays offer tremendous performance gains over standard HDD 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 any server system 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.

A simple glance at Hybrid RAID functionality does not readily show its common use cases, which include creating simple mirrors in workstations through to high-performance read-intensive applications in the small to medium business arena. Hybrid RAID is also used extensively in the data center to provide greater capacity in storage servers while providing fast boot for those servers. To learn more about Hybrid RAID, visit www.adaptec.com/en-us/solutions/hybrid-RAID/.

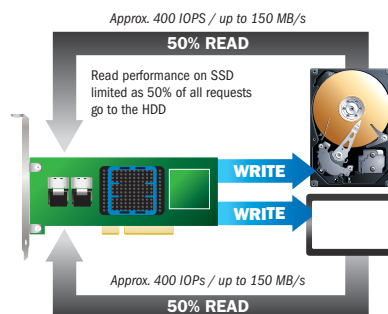


Figure 1 • Hybrid RAID Deployment

Who should use RAID?

Any server or high-end workstation, and any computer system where constant uptime is required, is a suitable candidate for RAID.

At some point in the life of a server, at least one drive will fail. Without some form of RAID protection, a failed drive's data would have to be restored from backups, likely at the loss of some data and a considerable amount of time.

With a RAID controller in the system, a failed drive can simply be replaced and the RAID controller will automatically rebuild the missing data from the rest of the drives onto the newly-inserted drive. This means that your system can survive a drive failure without the complex and long-winded task of restoring data from backups.

Choosing the right RAID level

There are several different RAID configurations, called "levels," such as RAID 0, RAID 1, RAID 10, and RAID 5. While there is little difference in their names, there are big differences in their characteristics and where/when they should be used.

The factors to consider when choosing the right RAID level include:

- Capacity
- Performance
- Redundancy (reliability/safety)
- Price

There is no one-size-fits all approach to RAID because focus on one factor typically comes at the expense of another. Some RAID levels designate drives to be used for redundancy, which means they can't be used for capacity. Other RAID levels focus on performance but not on redundancy. A large, fast, highly-redundant array will be expensive. Conversely, a small, average-speed redundant array won't cost much, but will not be anywhere near as fast as the previous expensive array.

With that in mind, here is a look at the different RAID levels and how they may meet your requirements.

RAID 0 (Striping)

In RAID 0, all drives are combined into one logical disk (Figure 2). This configuration offers low cost and maximum performance, but no data protection — a single drive failure results in total data loss.

As such, RAID 0 is not recommended. As SSDs become more affordable and grow in capacity, RAID 0 has declined in popularity. The benefits of fast read/write access are far outweighed by the threat of losing all data in the event of a drive failure.

Usage: Suited only for situations where data isn't mission-critical, such as video/audio post-production, multimedia imaging, CAD, data logging, etc. where it's OK to lose a complete drive because the data can be quickly re-copied from the source. Generally speaking, RAID 0 is not recommended.

Pros:

- Fast and inexpensive.
- All drive capacity is usable.
- Quick to set up. Multiple HDDs sharing the data load make it the fastest of all arrays.

Cons:

- RAID 0 provides no data protection at all.
- If one drive fails, all data will be lost with no chance of recovery.




Figure 2 • RAID 0 Array

RAID 1 (Mirroring)

RAID 1 maintains duplicate sets of all data on two separate drives while showing just one set of data as a logical disk (Figure 3). RAID 1 is about protection, not performance or capacity.

Since each drive holds copies of the same data, the usable capacity is 50% of the available drives in the RAID set.

Usage: Generally only used in cases where there is not a large capacity requirement, but the user wants to make sure the data is 100% recoverable in the case of a drive failure, such as accounting systems, video editing, gaming etc.

Pros:

- Highly redundant — each drive is a copy of the other.
- If one drive fails, the system continues as normal with no data loss.

Cons:

- Capacity is limited to 50% of the available drives, and performance is not much better than a single drive.

Note: With the advent of large-capacity SATA HDDs, it is possible to achieve an approximately 8 TB RAID 1 array using two 8 TB HDDs. While this may give sufficient capacity for many small business servers, performance will still be limited by the fact that it only has two spindles operating within the array. Therefore it is recommended to move to RAID arrays that utilize more spinning media when such capacities are required.




Figure 3 • RAID 1 Array

RAID 1E (Striped Mirroring)

RAID 1E combines data striping from RAID 0 with data mirroring from RAID 1 while offering more performance than RAID 1 (Figure 4). Data written in a stripe on one drive is mirrored to a stripe on the next drive in the array.

As in RAID 1, usable drive capacity in RAID 1E is 50% of the total available capacity of all drives in the RAID set.

Usage: Small servers, high-end workstations, and other environments with no large capacity requirements, but where the user wants to make sure the data is 100% recoverable in the case of a drive failure.

Pros:

- Redundant with better performance and capacity than RAID 1. In effect, RAID 1E is a mirror of an odd number of drives.

Cons:

- Cost is high because only half the capacity of the physical drives is available.




Figure 4 • RAID 1E Array

Note: RAID 1E is best suited for systems with three drives. For scenarios with four or more drives, RAID 10 is recommended.

RAID 5 (Striping with Parity)

As the most common and best “all-round” RAID level, RAID 5 stripes data blocks across all drives in an array (at least 3 to a maximum of 32), and also distributes parity data across all drives (Figure 5). In the event of a single drive failure, the system reads the parity data from the working drives to rebuild the data blocks that were lost.

RAID 5 read performance is comparable to that of RAID 0, but there is a penalty for writes since the system must write both the data block and the parity data before the operation is complete.

The RAID parity requires one drive capacity per RAID set, so usable capacity will always be one drive less than the total number of drives in the configuration.

Usage: Often used in file servers, general storage servers, backup servers, streaming data, and other environments that call for good performance but best value for the money. Not suited to database applications due to poor random write performance.

Pros:

- Good value and good all-around performance.

Cons:

- One drive capacity is lost to parity.
- Can only survive a single drive failure at any one time.
- If two drives fail at once, all data is lost.

Note: It is strongly recommended to have a hot spare set up with the RAID 5 to reduce exposure to multiple drive failures.

Note: While SSDs are becoming cheaper, and their improved performance over HDDs makes it seem possible to use them in RAID 5 arrays for database applications, the general nature of small random writes in RAID 5 still means that this RAID level should not be used in a system with a large number of small, random writes. A non-parity array such as RAID 10 should be used instead.




Figure 5 • RAID 5 Array

RAID 6 (Striping with Dual Parity)

In RAID 6, data is striped across several drives and dual parity is used to store and recover data (Figure 6). It is similar to RAID 5 in performance and capacity capabilities, but the second parity scheme is distributed across different drives and therefore offers extremely high fault tolerance and the ability to withstand the simultaneous failure of two drives in an array.

RAID 6 requires a minimum of 4 drives and a maximum of 32 drives to be implemented. Usable capacity is always two less than the number of available drives in the RAID set.

Usage: Similar to RAID 5, including file servers, general storage servers, backup servers, etc. Poor random write performance makes RAID 6 unsuitable for database applications.

Pros:

- Reasonable value for money with good all-round performance.
- Can survive two drives failing at the same time, or one drive failing and then a second drive failing during the data rebuild.

Cons:

- More expensive than RAID 5 due to the loss of two drive capacity to parity.
- Slightly slower than RAID 5 in most applications.




Figure 6 • RAID 6 Array

RAID 10 (Striping and Mirroring)

RAID 10 (sometimes referred to as RAID 1+0) combines RAID 1 and RAID 0 to offer multiple sets of mirrors striped together (Figure 7 and Figure 8). RAID 10 offers very good performance with good data protection and no parity calculations.

RAID 10 requires a minimum of four drives, and usable capacity is 50% of available drives. It should be noted, however, that RAID 10 can use more than four drives in multiples of two. Each mirror in RAID 10 is called a “leg” of the array. A RAID 10 array using, say, eight drives (four “legs,” with four drives as capacity) will offer extreme performance in both spinning media and SSD environments as there are many more drives splitting the reads and writes into smaller chunks across each drive.

Usage: Ideal for database servers and any environment with many small random data writes.

Pros:

- Fast and redundant.

Cons:

- Expensive because it requires four drives to get the capacity of two.
- Not suited to large capacities due to cost restrictions.
- Not as fast as RAID 5 in most streaming environments.




Figure 7 • RAID 10 Array – 2 Legs

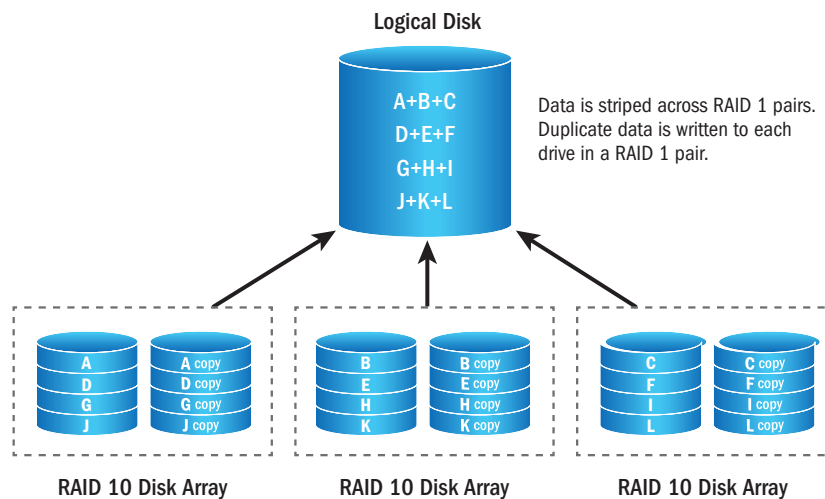


Figure 8 • RAID 10 Array – 3 Legs

RAID 50 (Striping with Parity)

RAID 50 (sometimes referred to as RAID 5+0) combines multiple RAID 5 sets (striping with parity) with RAID 0 (striping) (Figure 9 and Figure 10). The benefits of RAID 5 are gained while the spanned RAID 0 allows the incorporation of many more drives into a single logical disk. Up to one drive in each sub-array may fail without loss of data. Also, rebuild times are substantially less than a single large RAID 5 array.

A RAID 50 configuration can accommodate 6 or more drives, but should only be used with configurations of more than 16 drives. The usable capacity of RAID 50 is 67%-94%, depending on the number of data drives in the RAID set.

It should be noted that you can have more than two legs in a RAID 50. For example, with 24 drives you could have a RAID 50 of two legs of 12 drives each, or a RAID 50 of three legs of eight drives each. The first of these two arrays would offer greater capacity as only two drives are lost to parity, but the second array would have greater performance and much quicker rebuild times as only the drives in the leg with the failed drive are involved in the rebuild function of the entire array.

Usage: Good configuration for cases where many drives need to be in a single array but capacity is too large for RAID 10, such as in very large capacity servers.

Pros:

- Reasonable value for the expense.
- Very good all-round performance, especially for streaming data, and very high capacity capabilities.

Cons:

- Requires a lot of drives.
- Capacity of one drive in each RAID 5 set is lost to parity.
- Slightly more expensive than RAID 5 due to this lost capacity.

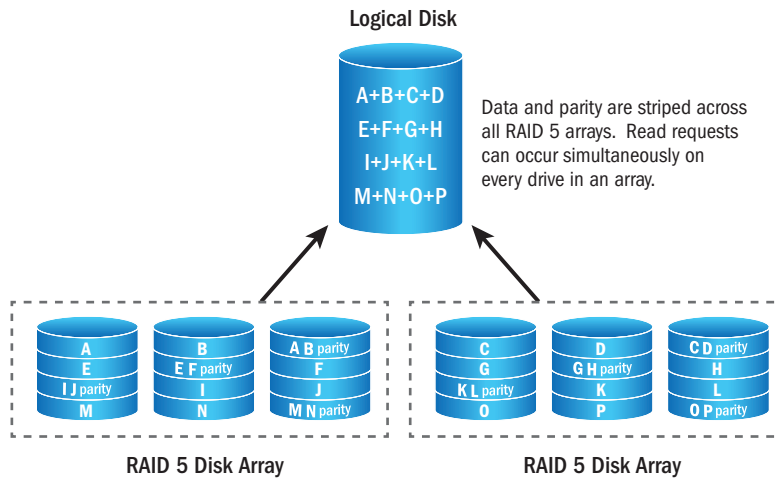


Figure 9 • RAID 50 Array – 2 Legs

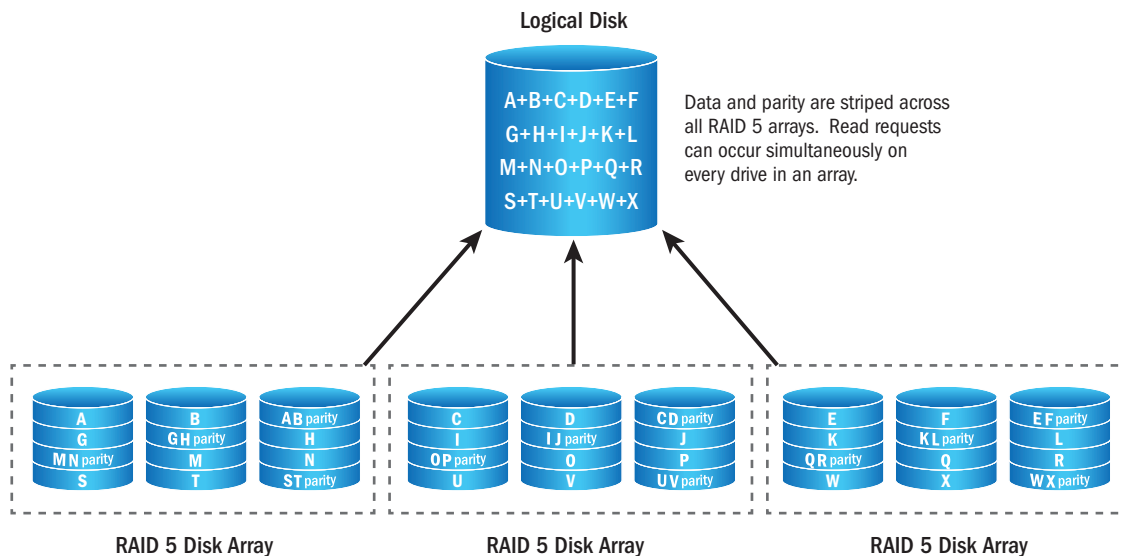


Figure 10 • RAID 50 Array – 3 Legs

RAID 60 (Striping with Dual Parity)

RAID 60 (sometimes referred to as RAID 6+0) combines multiple RAID 6 sets (striping with dual parity) with RAID 0 (striping) (Figure 11 and Figure 12). Dual parity allows the failure of two drives in each RAID 6 array while striping increases capacity and performance without adding drives to each RAID 6 array.

Like RAID 50, a RAID 60 configuration can accommodate 8 or more drives, but should only be used with configurations of more than 16 drives. The usable capacity of RAID 60 is between 50%-88%, depending on the number of data drives in the RAID set.

Note that all of the above multiple-leg configurations that are possible with RAID 10 and RAID 50 are also possible with RAID 60. With 36 drives, for example, you can have a RAID 60 comprising two legs of 18 drives each, or a RAID 60 of three legs with 12 drives in each.

Usage: RAID 60 is similar to RAID 50 but offers more redundancy, making it good for very large capacity servers, especially those that will not be backed up (i.e. video surveillance servers handling large numbers of cameras).

Pros:

Can sustain two drive failures per RAID 6 array within the set, so it is very safe.

Very large and reasonable value for money, considering this RAID level won't be used unless there are a large number of drives.

Cons:

Requires a lot of drives.

Slightly more expensive than RAID 50 due to losing more drives to parity calculations.

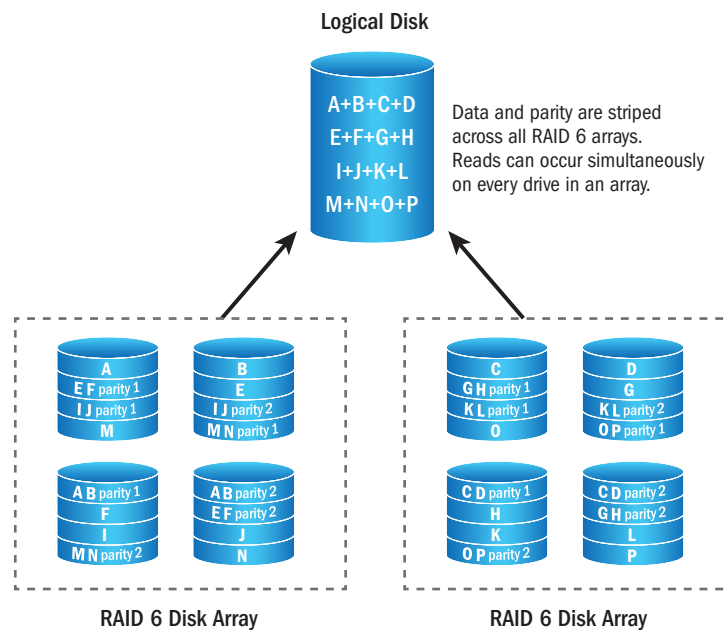


Figure 11 • RAID 60 Array – 2 Legs

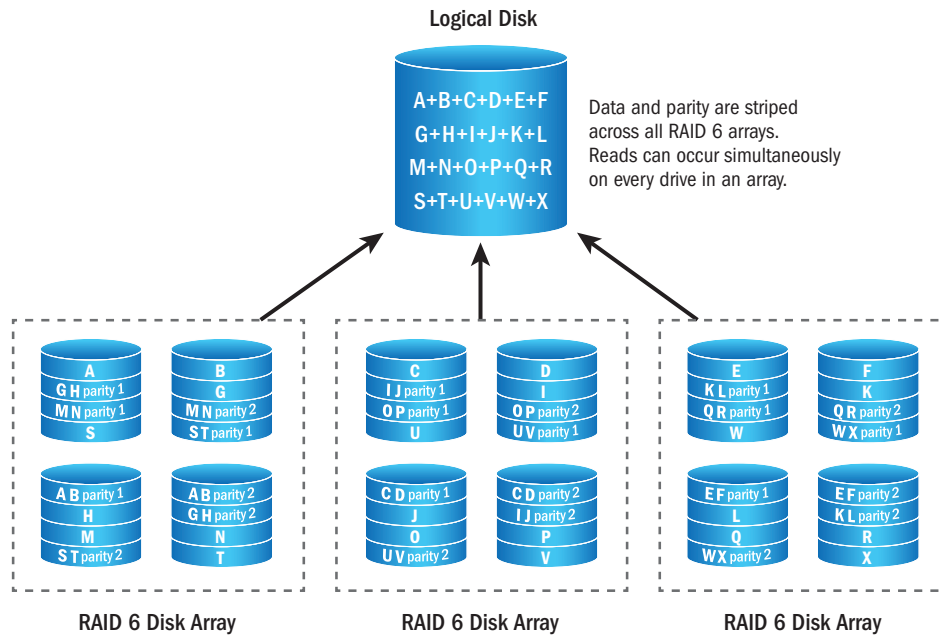


Figure 12 • RAID 60 Array – 3 Legs

When to use which RAID level

We can classify data into two basic types: random and streaming. As indicated previously, there are two general types of RAID arrays: non-parity (RAID 1, 10) and parity (RAID 5, 6, 50, 60).

Random data is generally small in nature (i.e., small blocks), with a large number of small reads and writes making up the data pattern. This is typified by database-type data.

Streaming data is large in nature, and is characterized by such data types as video, images, general large files.

While it is not possible to accurately determine all of a server's data usage, and servers often change their usage patterns over time, the general rule of thumb is that random data is best suited to non-parity RAID, while streaming data works best and is most cost-effective on parity RAID.

Note that it is possible to set up both RAID types on the same controller, and even possible to set up the same RAID types on the same set of drives. So if, for example, you have eight 2 TB drives, you can make a RAID 10 of 1 TB for your database-type data, and a RAID 5 of the capacity that is left on the drives for your general and/or streaming type data (approximately 12 TB). Having these two different arrays spanning the same drives will not impact performance, but your data will benefit in performance from being situated on the right RAID level.

Drive size performance

While HDDs are becoming larger, they are not getting any faster — a 1 TB HDD and a 6 TB HDD from the same product family will have the same performance characteristics. This has an impact when building/rebuilding arrays as it can take a long time to write all of the missing data to the new replacement drive.

Conversely, SSDs are often faster in larger capacities, so an 80 GB SSD and an 800 GB SSD from the same product family will have quite different performance characteristics. This should be checked carefully with the product specifications from the drive vendor to make sure you are getting the performance you think you are getting from your drives.

With HDDs it is generally better to create an array with more, rather than fewer, drives. A RAID 5 of three 6 TB HDDs (12 TB capacity) will not have the same performance as a RAID 5 array made from five 3 TB HDDs (12 TB capacity).

With SSDs, however, it is advisable to achieve the capacity required from as few as drives possible by using larger capacity SSDs. These will have higher throughput than their smaller counterparts and will yield better system performance.

Size of array vs size of drives

It is a little-known fact that you do not need to use all of your drive capacity when creating a RAID array. When, for example, creating the RAID array in the controller BIOS, the controller will show you the maximum possible size the array can be based on the drives chosen to make up the array.

During the creation process, you can change the size of the array to a lesser size. The unused space on the drives will be available for creating additional RAID arrays.

A good example of this would be when creating a large server and keeping the operating system and data on separate RAID arrays. Typically you would make a RAID 10 of, say, 200 GB for your OS installation spread across all drives in the server. This would use a minimal amount of capacity from each drive. You can then create a RAID 5 for your general data across the unused space on the drives.

This has an added benefit of getting around drive size limitations for boot arrays on non-UEFI servers as the OS will believe it is only dealing with a 200 GB drive when installing the operating system.

Rebuild times and large RAID arrays

The more drives in the array, and the larger the HDDs in the array, the longer the rebuild time when a drive fails and is replaced or a hot-spare kicks in. While it is possible to have

32 drives in a RAID 5 array, it becomes somewhat impractical to do this with large spinning media.

For example, a RAID 5 made of 32 6-TB drives (186 TB) will have very poor build and rebuild times due to the size, speed and number of drives. In this scenario, it would be advisable to build a RAID 50 with two legs from those drives (180 TB capacity). When a drive fails and is replaced, only 16 of the drives (15 existing plus the new drive) will be involved in the rebuild. This will improve rebuild performance and reduce system performance impact during the rebuild process.

Note, however, that no matter what you do, when it comes to rebuilding arrays with 6 TB+ SATA drives, rebuild times will increase beyond 24 hours in an absolutely perfect environment (no load on server). In a real-world environment with a heavily-loaded system, the rebuild times will be even longer.

Of course, rebuild times on SSD arrays are dramatically quicker due to the fact that the drives are smaller and the write speed of the SSDs are much faster than their spinning media counterparts.

Default RAID settings


When creating a RAID array in the BIOS or management software, you will be presented with defaults that the controller proposes for the RAID settings. The most important of these is the "stripe size." While there is much science, math and general knowledge involved in working out what is the best stripe size for your array, in the vast majority of cases the defaults work best, so use the 256kb stripe size as suggested by the controller.

SSDs and read/write cache

In an SSD-only RAID array, disabling the read and write cache will improve performance in a vast majority of cases. However, you may need test whether enabling read and write cache will improve performance even further. Note that it is possible to disable and enable read and write cache on the on the fly without affecting or reconfiguring the array, or restarting the server, so testing both configurations is recommended.

About Microsemi® Adaptec® RAID

While many companies offer RAID, not all RAID implementations are created equal. With nearly 35 years of protecting, accelerating, and conditioning data as it moves through the I/O path, Microsemi's Adaptec solution offers the most robust RAID data protection available today, including support for RAID levels: 0, 1, 1E, 5, 6, 10, 50, and 60, plus Hybrid RAID 1 and 10.



Summary

Table 1 • RAID Level Comparison (RAID 0, 1, 1E)

Features	RAID 0	RAID 1	RAID 1E
Minimum # Drives	2	2	3
Data Protection	No Protection	Single-drive failure	Single-drive failure
Read Performance	High	Medium	Medium
Write Performance	High	Medium	Medium
Read Performance (degraded)	N/A	Medium	Medium
Write Performance (degraded)	N/A	Medium	Medium
Capacity Utilization	100%	50%	50%
Typical Usage	<p>Non-mission-critical data, such as video/audio post-production, multimedia imaging, CAD, data logging, etc. where it's OK to lose a complete drive because the data can be quickly re-copied from the source.</p> <p>GENERALLY SPEAKING, RAID 0 IS NOT RECOMMENDED.</p>	<p>Cases where there is not a large capacity requirement, but the user wants to make sure the data is 100% recoverable in the case of a drive failure, such as accounting systems, video editing, gaming etc.</p>	<p>Small servers, high-end workstations, and other environments where no large capacity requirements, but where the user wants to make sure the data is 100% recoverable in the case of a drive failure.</p>

Table 2 • RAID Level Comparison (RAID 6, 10, 50, 60)


Features	RAID 6	RAID 10	RAID 50	RAID 60
Minimum # Drives	4	4	6	8
Data Protection	Two-drive failure	Up to one drive failure in each sub-array	Up to one drive failure in each sub-array	Up to two drive failure in each sub-array
Read Performance	High	High	High	High
Write Performance	Low	High	Medium	Medium
Read Performance	Low	High	Medium	Medium
Write Performance	Low	High	Medium	Low
Capacity Utilization	50% - 88%	50%	67% - 94%	67% - 94%

Table 2 • RAID Level Comparison (RAID 6, 10, 50, 60)

Features	RAID 6	RAID 10	RAID 50	RAID 60
Typical Usage	Similar to RAID 5, including file servers, general storage servers, backup servers, etc. Poor random write performance makes RAID 6 unsuitable for database applications.	Ideal for database servers and any environment with many small random data writes.	Good configuration for cases where many drives need to be in a single array but capacity is too large for RAID 10, such as in very large capacity servers.	RAID 60 is similar to RAID 50 but offers more redundancy, making it good for very large capacity servers, especially those that will not be backed up (i.e. video surveillance servers handling large numbers of cameras).
Pros	Reasonable value for money with good all-round performance. Can survive two drives failing at the same time, or one drive failing and then a second drive failing during the data rebuild.	Fast and redundant.	Reasonable value for the expense. Very good all-round performance, especially for streaming data, and very high capacity capabilities.	Can sustain two drive failures per RAID 6 array within the set, so it is very safe. Reasonable value for money, considering this RAID level won't be used unless there are a large number of drives.
Cons	More expensive than RAID 5 due to the loss of two drive capacity to parity. Slightly slower than RAID 5 in most applications.	Expensive as it requires four drives to get the capacity of two. Not suited to large capacities due to cost restrictions. Not as fast as RAID 5 in streaming environments.	Requires a lot of drives. Capacity of one drive in each RAID 5 set is lost to parity. Slightly more expensive than RAID 5 due to this lost capacity.	Requires a lot of drives. Slightly more expensive than RAID 50 due to losing more drives to parity calculations.

Table 3 • Types of RAID

Types of RAID	Software-Based	Motherboard-Based	Adapter-Based
Description	Included in the OS, such as Windows®, and Linux. All RAID functions are handled by the host CPU which can severely tax its ability to perform other computations.	Processor-intensive RAID operations are off-loaded from the host CPU to a RAID processor integrated into the motherboard.	Processor-intensive RAID operations are off-loaded from the host CPU to an external PCIe adapter. Battery-back write back cache can dramatically increase performance without adding risk of data loss.
Typical Usage	Best used for large block applications such as data warehousing or video streaming. Also where servers have the available CPU cycles to manage the I/O intensive operations certain RAID levels require.	Inexpensive.	Best used for small block applications such as transaction oriented databases and web servers.
Pros	Lower cost due to lack of RAID-dedicated hardware.	Lower cost than adapter-based RAID.	Offloads RAID tasks from the host system, yielding better performance than software RAID. Controller cards can be easily swapped out for replacement and upgrades. Data can be backed up to prevent loss in a power failure.
Cons	Lower RAID performance as CPU also powers the OS and applications.	No ability to upgrade or replace the RAID processor in the event of hardware failure. May only support a few RAID levels.	More expensive than software and integrated RAID.



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

© 2016 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

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.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.