RAID: The Basics

University Training Module

March 2016

RAID: The Basics

Introduction

Servers require bigger, faster, and safer storage than single-disk technology can provide. How do we increase storage space, increase read and write speeds, and protect data from disk drive failure?

RAID is the answer. With RAID, we can combine multiple disk drives to create large, safe virtual (or logical) drives, which are also far faster than any single disk.

RAID allows for many configurations of disk drives, but RAID's key principles of redundancy, performance, and capacity allow system administrators to build servers whose storage meets the requirements of their specific organization.

In this module, you'll learn more about RAID and how you can use it to develop a storage system that meets an organization's needs at a cost point it can afford.

At the completion of this module, you should be able to:

- Explain what RAID is and why it's important
- Understand some basic RAID concepts
- Identify the nine different RAID levels and understand how they differ
- Feel comfortable choosing between RAID levels to meet different business requirements

What is RAID?

RAID stands for Redundant Array of Independent Disks and is a method of treating several individual disk drives as one logical unit (array or logical drive) to improve security, reliability, and read and write performance.

How does RAID work?

RAID increases data protection and performance by duplicating and/or spreading data over multiple disks.

Data can be mirrored (duplicated on multiple disks), or striped (spread over multiple disks), or both, depending on the RAID level.

Nine RAID levels (0, 1, 1E, 5, 5EE, 6, 10, 50, 60) provide varying levels of performance and reliability based on the number of disk drives in the array. (RAID levels are covered later in this module.)

Why is RAID important?

RAID protects against data loss and provides real-time data recovery when a disk drive fails, reducing system downtime.

Multiple drives working together also offer better read and write performance than a single disk drive or group of independent disk drives.

Key RAID Concepts

Before continuing, make sure you know these important RAID concepts:

Fault Tolerance. The ability of a system to continue to perform its functions even when one or more of its disk drives have failed.

Striping. Spreads data evenly over multiple disk drives. Offers no data protection, but increases performance. Requires at least two disk drives.

Mirroring. Duplicates data from a primary disk drive to a secondary disk drive to safeguard against data loss. Requires two disk drives.

Mirroring and Striping. Stripes data across mirror sets to increase performance. Requires at least four disk drives.

Parity. A calculation based on stored data that's used to provide redundancy in an array. If a disk drive in an array fails, parity data is combined with the data on the remaining disk drives to recreate the missing information.

Drive Segment. Storage space on a disk drive that is used to create an array. A drive segment can comprise all or only part of the space on a disk drive. A single disk drive can be part of one or more arrays, but a single drive segment can be part of only one array at a time.

Logical Drive. Another term for an array—a group of disk drives (or drive segments) treated as a single unit. Also known as a logical device or a container. Adaptec Storage Manager, our storage management software, refers to arrays as 'logical drives'.

Hot Spare. A spare disk drive which is used as an automatic replacement for any failed disk drive in a RAID array. A hot spare can be shared by multiple arrays and must be as large as the smallest member of the array.

Dedicated Hot Spare. A spare disk drive that's assigned to a specific array. Must be as large as the smallest member of the array.

Copyback Hot Spare. A hot spare with special features: When a failed disk drive is replaced, data is automatically copied from the hot spare back to the replaced drive.

Distributed Spare. Unlike a hot spare, a distributed spare is striped evenly across the disk drives with the stored data and parity data, and can't be shared with other arrays. A distributed spare improves the speed at which the array is rebuilt following a disk drive failure.

Hot Swap. The ability to replace a failed member of a RAID array with a functioning disk drive without bringing down the server or interrupting transactions that involve the other disk drives in the array.

RAID Levels

RAID 0 (Non-redundant Arrays)

An array with RAID 0 includes two or more disk drives and provides striping (see page 2). Compared to an equal-sized group of independent disks, a RAID 0 array provides improved performance.

Note: RAID 0 arrays do not maintain redundant data, so they cannot recover from single disk drive failures and offer no data protection.

Drive segment size is limited to the size of the smallest disk drive in the array. For instance, an array with two 250 GB disk drives and two 400 GB disk drives can create a RAID 0 drive segment of 250 GB, for a total of 1000 GB for the volume, as shown in this figure.

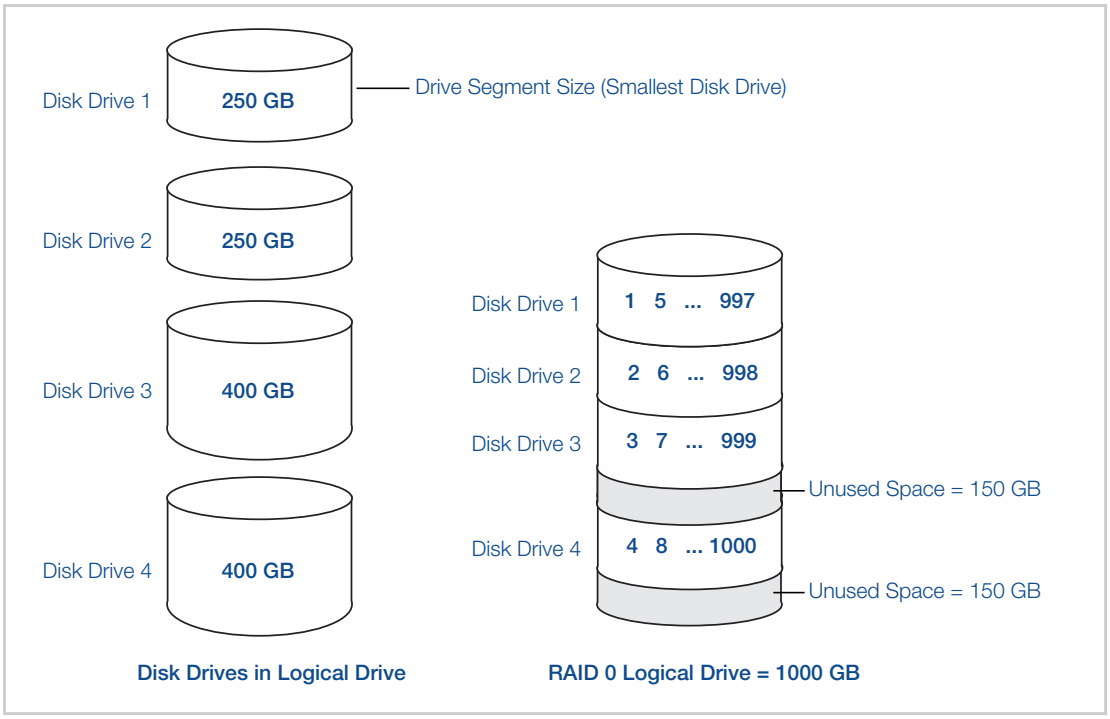

RAID 1

A RAID 1 array is built from two disk drives, where one disk drive is a mirror of the other.

Compared to independent disk drives, RAID 1 arrays provide improved performance, with a slightly faster read rate and an equal write rate of single disks. RAID 1 arrays protect your data against single disk drive failures. However, capacity is only 50 percent of two independent disk drives.

If the RAID 1 array is built from different-sized disk drives, drive segment size is the size of the smaller disk drive, as shown in this figure.

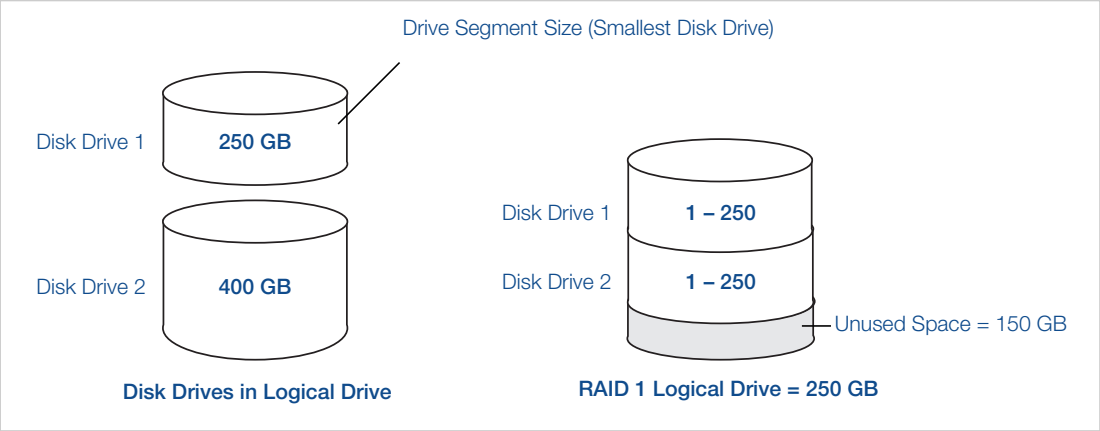

RAID 1E (RAID 1 Enhanced)

A RAID 1E array—also referred to as a striped mirror—is similar to a RAID 1 array except that data is both mirrored and striped (see page 2), and more disk drives can be included. A RAID 1E array can be built from three or more disk drives.

RAID 1E arrays protect your data against single disk drive failures.

In this figure, the green numbers represent the striped data, and the blue numbers represent the mirrored data stripes.

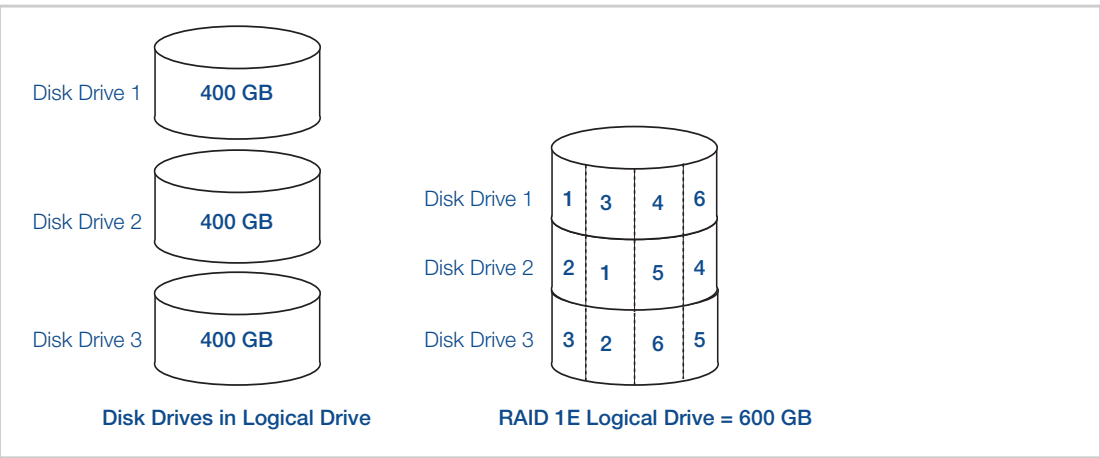

RAID 10

A RAID 10 array uses a minimum of 4 drives to stripe data across mirror sets. Data in a RAID 10 array is both striped and mirrored (see page 2). Mirroring provides data protection and striping improves performance. RAID 10 arrays protect your data against single disk drive failures.

Drive segment size is limited to the size of the smallest disk drive in the array. For instance, an array with two 250 GB disk drives and two 400 GB disk drives can create two mirrored drive segments of 250 GB, for a total of 500 GB for the array, as shown in this figure.

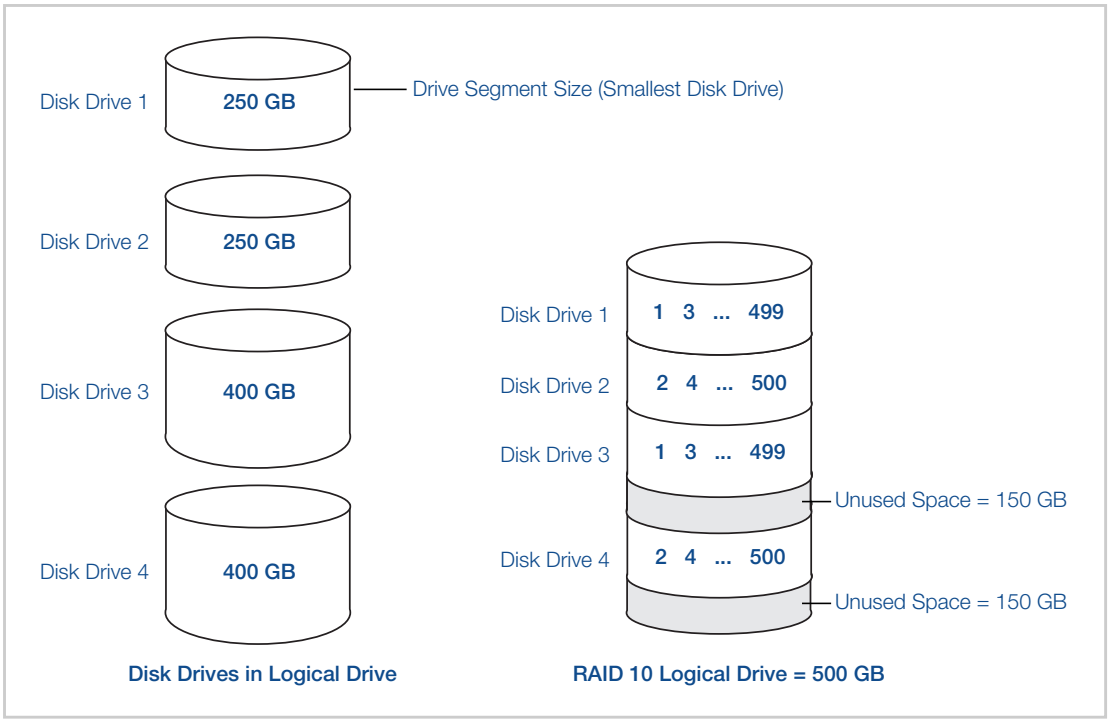

RAID 5

A RAID 5 array is built from a minimum of three disk drives, and uses data striping and parity data (see page 2) to provide redundancy. Parity data provides data protection, and striping improves performance. RAID 5 arrays protect your data against single disk drive failures.

In RAID 5 arrays, parity data (represented by Ps in the next figure) is striped evenly across the disk drives with the stored data.

Drive segment size is limited to the size of the smallest disk drive in the array. For instance, an array with two 250 GB disk drives and two 400 GB disk drives can contain 750 GB of stored data and 250 GB of parity data, as shown in this figure.

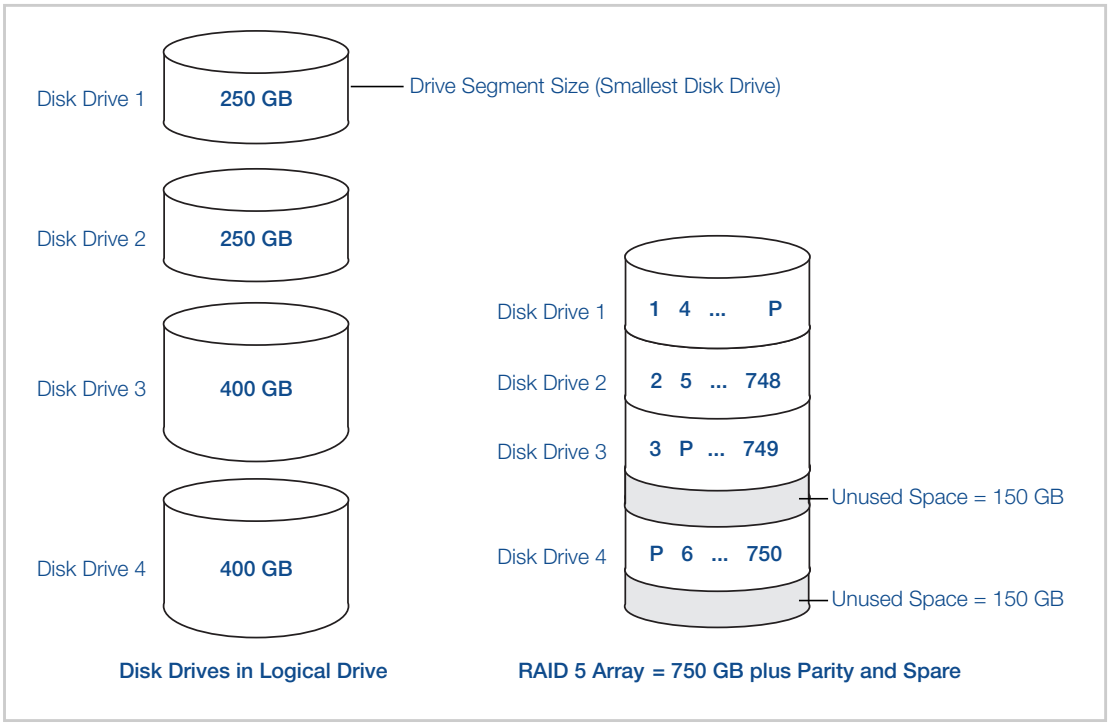

RAID 5EE (Hot Space)

A RAID 5EE array—also referred to as a hot space—is similar to a RAID 5 array except that it includes a distributed spare (see page 2) and must be built from a minimum of four disk drives.

A RAID 5EE array protects your data and increases read and write speeds. RAID 5EE arrays protect your data against single disk drive failures. However, capacity is reduced by two disk drives' worth of space, which is for parity data and spare data.

In this example, S represents the distributed spare, P represents the distributed parity data.

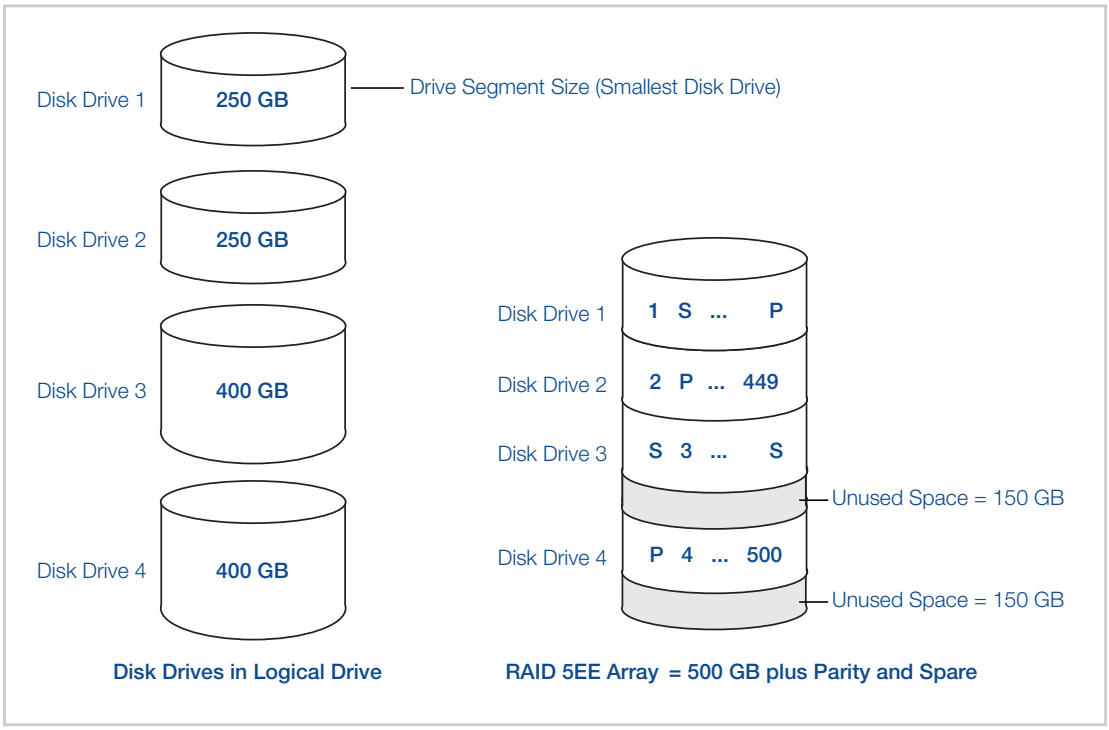

RAID 50

A RAID 50 array is built from at least six disk drives configured as two or more RAID 5 arrays, and stripes stored data and parity data across all disk drives in both RAID 5 arrays.

The parity data provides data protection, and striping improves performance. RAID 50 arrays protect your data against single disk drive failures. RAID 50 arrays also provide high data transfer speeds.

Drive segment size is limited to the size of the smallest disk drive in the array. For example, three 250 GB disk drives and three 400 GB disk drives comprise two equal-sized RAID 5 arrays with 500 GB of stored data and 250 GB of parity data. The RAID 50 array can therefore contain 1000 GB (2 x 500 GB) of stored data and 500 GB of parity data.

In this example, P represents the distributed parity data.

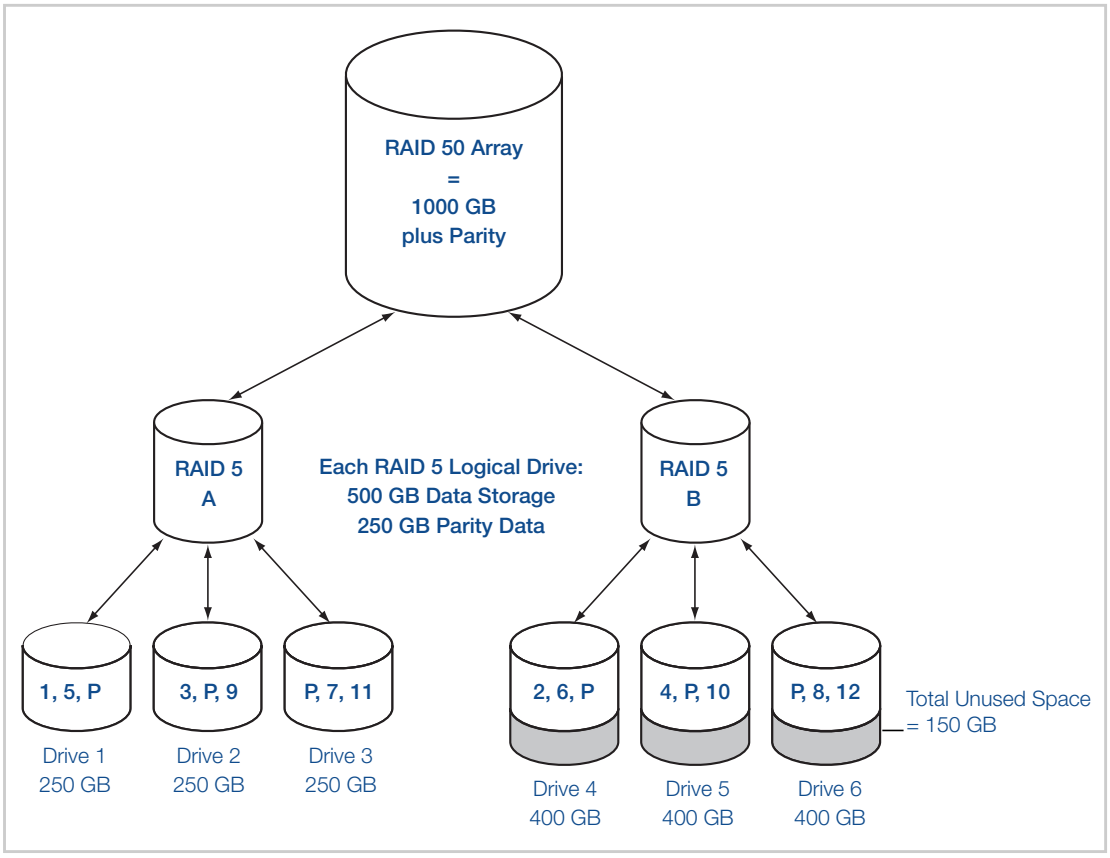

RAID 6 (Dual Drive Failure Protection)

A RAID 6 array—also referred to as dual drive failure protection—is similar to a RAID 5 array because it uses data striping and parity data to provide redundancy. However, RAID 6 arrays include two independent sets of parity data instead of one. Both sets of parity data are striped separately across all disk drives in the array.

RAID 6 arrays provide extra protection for your data because they can recover from two simultaneous disk drive failures. However, the extra parity calculation slows performance compared to RAID 5 arrays.

RAID 6 arrays must be built from at least four disk drives.

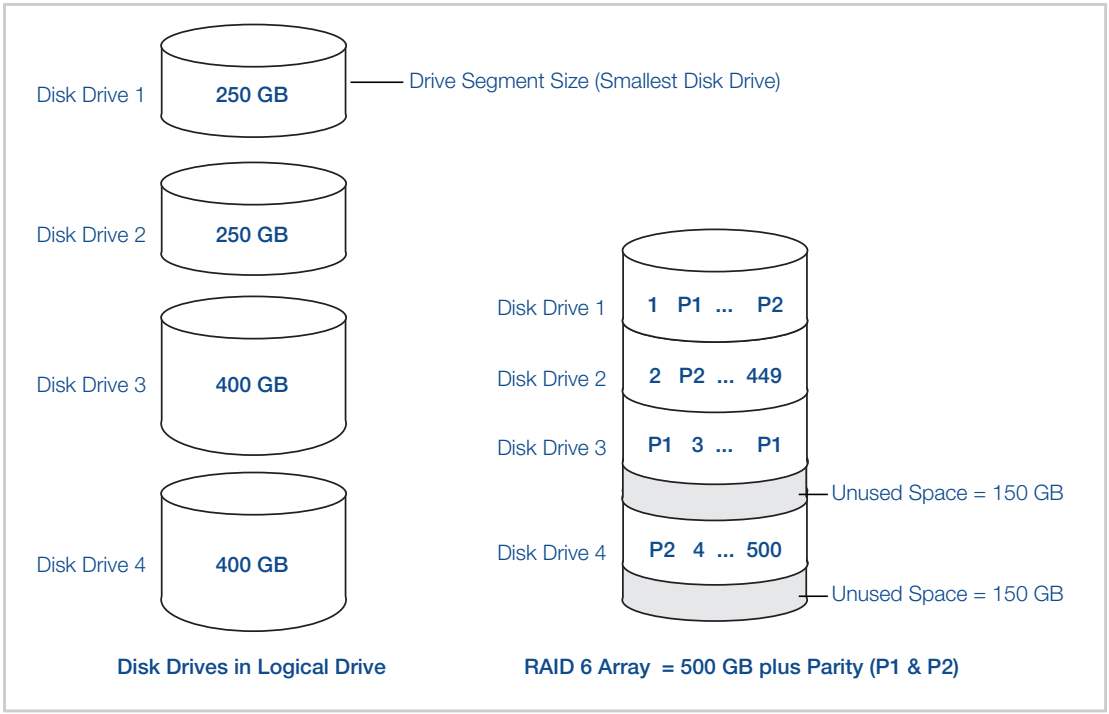

RAID 60 (Dual Drive Failure Protection)

Similar to a RAID 50 array (see page 8), a RAID 60 array—also referred to as dual drive failure protection— is built from at least eight disk drives configured as two or more RAID 6 arrays, and stripes stored data and two sets of parity data across all disk drives in both RAID 6 arrays.

Two sets of parity data provide enhanced data protection, and striping improves performance. RAID 60 arrays provide extra protection for your data because they can recover from two simultaneous disk drive failures. RAID 60 arrays also provide high data transfer speeds.

Working with RAID

How do I select the right RAID level?

This table helps you select the most appropriate RAID level for an array, based on the number of available disk drives and the requirements for performance and reliability.

Disk drive usage, read performance, and write performance depend on the number of drives in the array. In general, the more drives, the better the performance.

RAID Level	Redundancy	Disk Drive Usage	Read Performance	Write Performance	Built-in Hot Spare	Minimum Disk Drives	Benefit
RAID 0	No	100%	★★★	★★★	No	2	Highest performance
RAID 1	Yes	50%	★	★	No	2	Best data protection
RAID 1E	Yes	50%	★★	★★	No	3	Best data protection for an odd number of disk drives
RAID 10	Yes	50%	★★	★★	No	4	Highest performance with best data protection
RAID 5	Yes	67% – 94%	★★★	★	No	3	Best cost/performance balance for multi-disk drive environments
RAID 5EE	Yes	50% – 88%	★★★	★	Yes	4	The cost/performance balance of RAID 5 without setting aside a dedicated hot spare
RAID 50	Yes	67% – 94%	★★★	★	No	6	Increased capacity and performance for multi-array RAID 5 environments
RAID 6	Yes	50% – 88%	★★	★	No	4	Highest fault tolerance with the ability to survive two simultaneous disk drive failures
RAID 60	Yes	50% – 88%	★★	★	No	8	Highest performance with best data protection and the ability to survive two simultaneous disk drive failures

Can I change the RAID level of an existing array?

You can change the RAID level of an existing array from one RAID level to another, with some limitations. (Changing the RAID level is referred to as migrating.) Data on a migrating array remains available during the conversion process, but performance is significantly reduced until the migration is complete.

This table shows which existing arrays can be migrated to another RAID level:

Existing	Can be converted to
RAID 0	RAID 5 or 10
RAID 1	RAID 0 or 5 or 10
RAID 5	RAID 0 or 10 or 6
RAID 6	RAID 5
RAID 10	RAID 0 or 5
RAID 50	RAID 0 or 5 or 10 or 60

Note: Converting from some RAID levels may require assigning additional disks/segments to the new array to meet the requirements of the RAID level you want to use. Converting to RAID 6 or 60 can result in a reduction of usable capacity due to the additional parity used for RAID 6.

How large can a RAID array be?

In terms of maximum sizes for a RAID array, what is technically possible and what is recommended for real-world application are two different things. This table shows the realistic maximums recommended by Adaptec, compared to what RAID technology can support in a strictly academic sense.

	Maximum Recommended	Technically Possible
Maximum number of arrays per adapter	4	64
Maximum number of arrays that can be created on the same set of drives	4	4
Maximum array size	512 TB[a]	512TB
Maximum number of drives in RAID 0	0[b]	128
Maximum number of drives in RAID 5 array	8–12	32
Maximum number of drives in RAID 50 array	8–12 per leg	32 per leg
Maximum number of drives in RAID 6 array	8–12	32
Maximum number of drives in RAID 60 array	8–12 per leg	32 per leg
Overall maximum size for a RAID array	512TB	512TB

a. **Keep in mind that the larger the size, the longer it takes to build and rebuild the array.**
b. **RAID 0 is not recommended, as it provides no data protection.**

Can I make a RAID array larger if I need to?

You can expand an array to include additional storage space (disk segments), with some limitations. Any RAID 0, RAID 5, or RAID 6 array can be expanded to include the maximum recommended number of disk drives.

Caution: A RAID migration and a capacity expansion should not be done simultaneously.

What happens if a disk drive in a RAID array fails?

What happens after a disk drive failure depends on the RAID level assigned to the array and whether or not there was a hot spare (dedicated, distributed, or copyback—see page 2) available.

- Because RAID 0 arrays have no redundancy built in, a disk drive failure means that data is lost and is unrecoverable.
- In a RAID 1 array, which consists of two exact copies of the same data, the remaining disk drive can be used to rebuild the array with a new functioning disk drive.
- In all other RAID levels, a single disk drive failure causes the array to be rebuilt using a hot spare, assuming one is available.
- In all RAID levels except for RAID 6 and RAID 60, two simultaneous disk drive failures mean that data is lost and unrecoverable.

Is RAID all the protection I need for my data?

RAID is not all the protection you need for your data. RAID not a substitute for good backups. Adaptec strongly recommends that you maintain adequate backups as well as storing data on a RAID array.

For the most recent updates to our product line and for detailed information and specifications, please call, email or visit our website:

Toll-free: 800-713-4113

sales.support@microsemi.com

www.microsemi.com

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
email: sales.support@microsemi.com
www.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; 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, Calif., and has approximately 4,800 employees globally. Learn more at www.microsemi.com.