adaptec

Which RAID Level is Right for Me?

Contents

Data is the most valuable asset of any business today. Lost data means lost business. Even if you backup regularly, you need a fail-safe way to ensure that your data is protected and can be accessed without interruption in the event of an online disk failure. Adding RAID to your storage configurations is one of the most cost-effective ways to maintain both data protection and access.

While a number of companies offer RAID, not all RAID implementations are created equal. With over 24 years of SCSI development experience, only Adaptec offers the most robust RAID data protection available today, based on a hardened RAID code proven over years of use in demanding environments and resold by most of the top-tier computer manufacturers.

To choose the RAID level that's right for you, begin by considering the factors below. Each one of these factors becomes a trade-off for another:

- Cost of disk storage
- Data protection or data availability required (low, medium, high)
- Performance requirements (low, medium, high)

Cost boils down to the trade-off between disk capacity and added data availability or performance. For example, RAID 1/10 and small disk counts of RAID 6 are costly in terms of lost disk space (50%), but high in data availability.

Performance also depends on the access pattern (random/sequential, read/write, long/short) and the numbers of users. This white paper intends to give an overview on the performance and availability of various RAID levels in general and may not be accurate in all user scenarios.

RAID Level Descriptions:

RAID 0 (Striping):

Offers low cost and maximum performance, but offers no fault tolerance. A single disk failure results in TOTAL data loss. Businesses use RAID 0 mainly for tasks requiring fast access to a large capacity of temporary disk storage (such as video/audio post-production, multimedia imaging, CAD, data logging, etc.) where in case of a disk failure, the data can be easily reloaded without impacting the business. There are also no cost disadvantages as all storage is usable. RAID 0 usable capacity is 100% as all available drives are used.

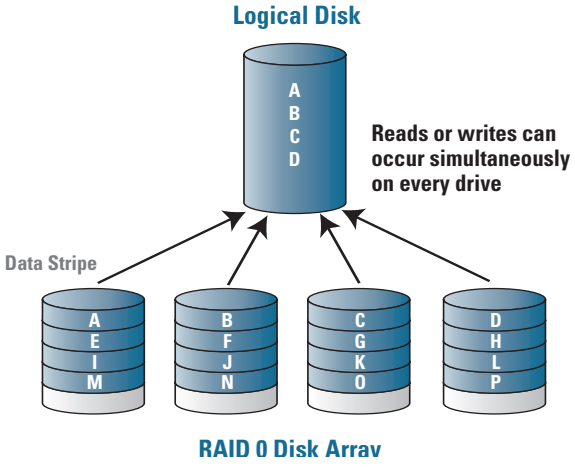

RAID 1 (Mirroring):

Provides cost-effective, high fault tolerance for configurations with two disk drives. RAID 1 refers to maintaining duplicate sets of all data on separate disk drives. It also provides the highest data availability since two complete copies of all information are maintained. There must be two disks in the configuration and there is a cost disadvantage as the usable capacity is half the number of available disks. RAID 1 offers data protection insurance for any environments where absolute data redundancy, availability and performance are key, and cost per usable gigabyte of capacity is a secondary consideration.

RAID 1 usable capacity is 50% of the available drives in the RAID set.

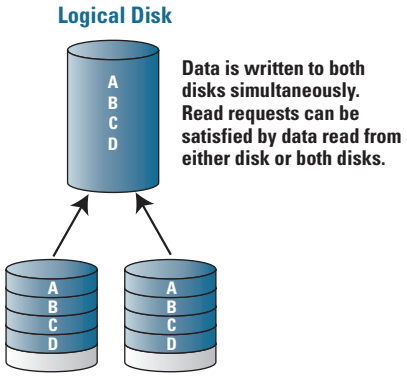

RAID 1 Disk Array

RAID 1E (Striped mirroring)

Combines data striping from RAID 0 with data mirroring from RAID 1. Data written in a stripe on one disk is mirrored to a stripe on the next drive in the array. The main advantage over RAID 1 is that RAID 1E arrays can be implemented using an odd number of disks.

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

Note: When using even numbers of disks it is always preferable to use RAID 10, which will allow multiple drive failures. With odd numbers of disks, however, RAID 1E supports only one drive failure.

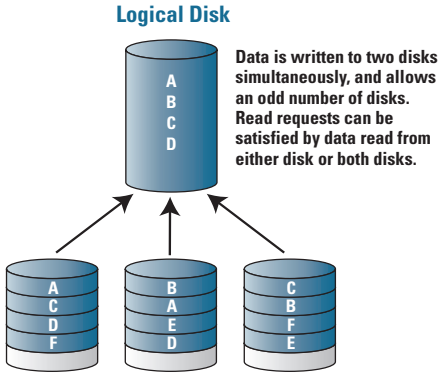

RAID 1E Disk Array

RAID 5 (Striping with parity):

Uses data striping in a technique designed to provide fault-tolerant data storage, but doesn't require duplication of data like RAID 1 and RAID 1E. Data is striped across all of the drives in the array, but for each stripe through the array (one stripe unit from each disk) one stripe unit is reserved to hold parity data calculated from the other stripe units in the same stripe. Read performance is therefore very good, but there is a penalty for writes, since the parity data has to be recalculated and written along with the new data. To avoid a single drive bottleneck, the parity data for consecutive stripes is interleaved with the data across all disks in the array.

RAID 5 has been the standard in server environments requiring fault tolerance. The RAID parity requires one disk drive per RAID set, so usable capacity will always be one disk drive less than the number of available disks in the configuration of available capacity — still better than RAID 1 which has only a 50% usable capacity.

RAID 5 requires a minimum of three disks and a maximum of 16 disks to be implemented. RAID 5 usable capacity is between 67% - 94%, depending on the number of data drives in the RAID set.

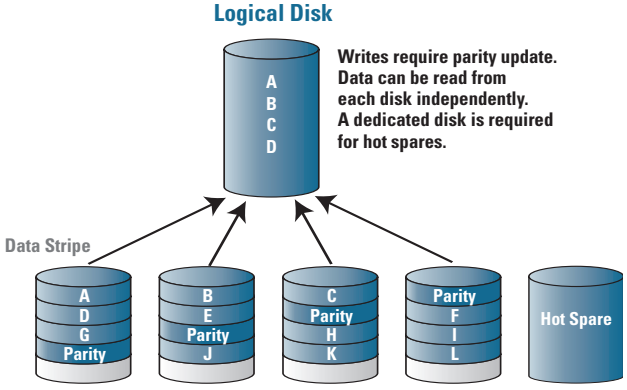

RAID 5 Disk Array

RAID 5EE (Hot Space):

Provides the protection of RAID 5 with higher I/Os per second by utilizing one more drive, with data efficiently distributed across the spare drive for improved I/O access.

RAID 5EE distributes the hot-spare drive space over the N+1 drives comprising the RAID-5 array plus standard hot-spare drive. This means that in normal operating mode the hot spare is an active participant in the array rather than spinning unused. In a normal RAID 5 array adding a hot-spare drive to a RAID 5 array protects data by reducing the time spent in the critical degraded state. However, this technique does not make maximum use of the hot-spare drive because it sits idle until a failure occurs. Often many years can elapse before the hot-spare drive is ever used. For small RAID 5 arrays in particular, having an extra disk to read from can provide significantly better read performance.

For example, going from a 4-drive RAID 5 array with a hot spare to a 5-drive RAID 5EE array will increase performance by roughly 25%.

One downside of RAID 5EE is that the hot-spare drive cannot be shared across multiple physical arrays as with standard RAID 5 plus hot-spare. This RAID 5 technique is more cost-efficient for multiple arrays because it allows a single hot-spare drive to provide coverage for multiple physical arrays. This configuration reduces the cost of using a hot-spare drive, but the downside is the inability to handle separate drive failures within different arrays. This RAID level can sustain a single drive failure.

RAID 5EE useable capacity is between 50% - 88%, depending on the number of data drives in the RAID set. RAID 5EE requires a minimum of four disks and a maximum of 16 disks to be implemented.

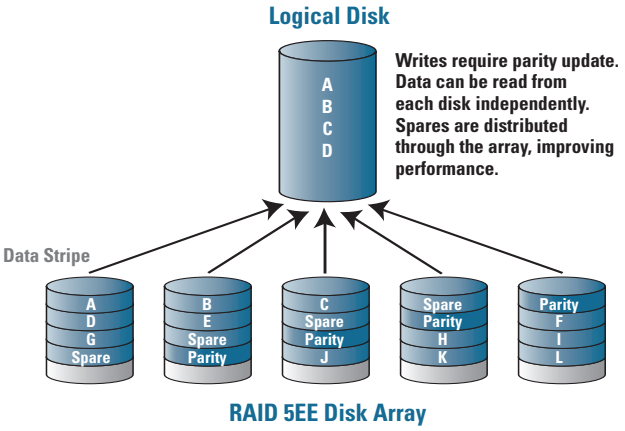

RAID 6 (Striping with dual parity):

Data is striped across several physical drives and dual parity is used to store and recover data. It tolerates the failure of two drives in an array, providing better fault tolerance than RAID 5. It also enables the use of more cost-effective, but less reliable, ATA and SATA disks for business critical data.

This RAID level is similar to RAID 5, but includes a second parity scheme that is distributed across different drives and therefore offers extremely high fault tolerance and drive-failure tolerance. RAID 6 can withstand a double disk failure.

RAID 6 requires a minimum of four disks and a maximum of 16 disks to be implemented. Usable capacity is always 2 less than the number of available disk drives in the RAID set.

Note: With less expensive, but less reliable SATA disk drives in a configuration that employs RAID 6, it is possible to achieve a higher level of availability than a Fibre Channel Array using RAID 5. This is because the second parity drive in the RAID 6 RAID set can withstand a second failure during a rebuild. In a RAID 5 set, the degraded state and/or the rebuilding time onto a hot spare is considered the window at which the RAID array is most vulnerable to data loss. During this time, if a second disk failure occurs, data is unrecoverable. With RAID 6 there are no windows of vulnerability as the second parity drive protects against this.

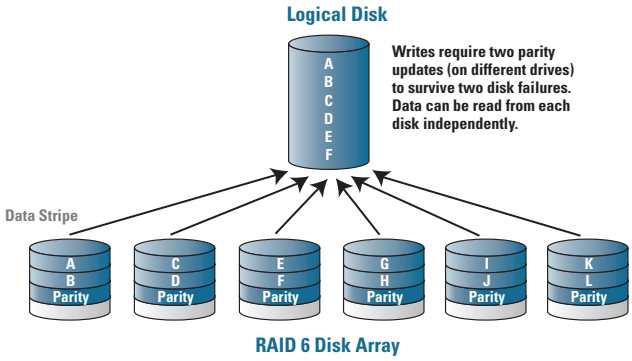

RAID 10 (Striped RAID 1 sets):

Combines RAID 0 striping and RAID 1 mirroring. This level provides the improved performance of striping while still providing the redundancy of mirroring.

RAID 10 is the result of forming a RAID 0 array from two or more RAID 1 arrays. This RAID level provides fault tolerance — up to one disk of each sub-array may fail without causing loss of data.

Usable capacity of RAID 10 is 50% of available disk drives.

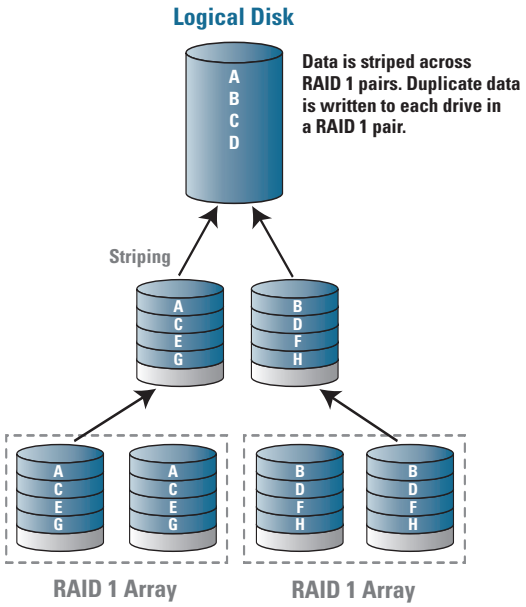

RAID 10 Disk Array

RAID 50 (Striped RAID 5 sets):

Combines multiple RAID 5 sets with RAID 0 (striping). Striping helps to increase capacity and performance without adding disks to each RAID 5 array (which will decrease data availability and could impact performance when running in a degraded mode).

RAID 50 comprises RAID 0 striping across lower-level RAID 5 arrays. The benefits of RAID 5 are gained while the spanned RAID 0 allows the incorporation of many more disks into a single logical drive. Up to one drive in each sub-array may fail without loss of data. Also, rebuild times are substantially less then a single large RAID 5 array.

Usable capacity of RAID 50 is between 67% - 94%, depending on the number of data drives in the RAID set.

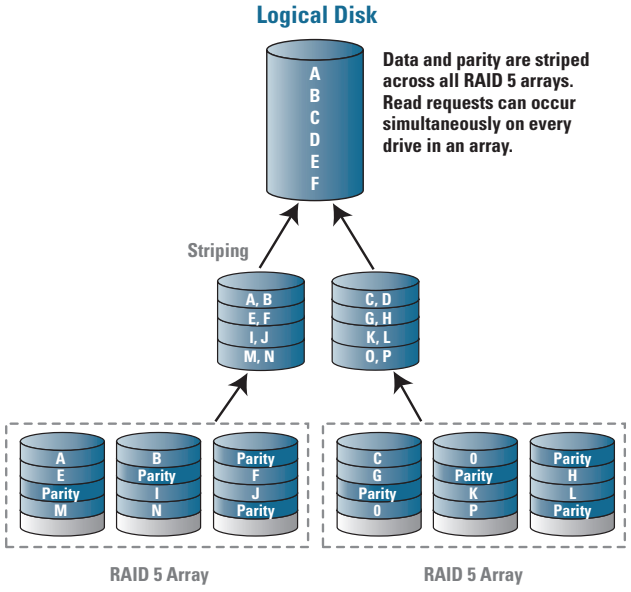

RAID 50 Disk Array

RAID 60 (Striped RAID 6 sets):

Combines multiple RAID 6 sets with RAID 0 (striping). Dual parity allows the failure of two disks in each RAID 6 array. Striping helps to increase capacity and performance without adding disks to each RAID 6 array (which would decrease data availability and could impact performance in degraded mode).

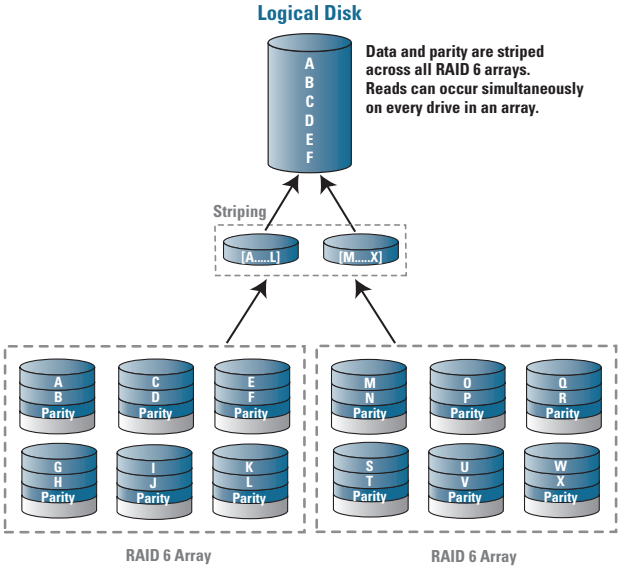

RAID 60 Disk Array

Which RAID Level is Right for Me?

RAID Level Comparison

Features	RAID 0	RAID 1	RAID 1E	RAID 5	RAID 5EE	RAID 6	RAID 10	RAID 50	RAID 60
Minimum # Drives	2	2	3	3	4	4	4	6	8
Data Protection	No Protection	Single-drive failure	Single-drive failure	Single-drive failure	Single-drive failure	Two-drive failure	Up to one disk failure in each sub-array	Up to one disk failure in each sub-array	Up to two disk failures in each sub-array
Read Performance	High	High	High	High	High	High	High	High	High
Write Performance	High	Medium	Medium	Low	Low	Low	Medium	Medium	Medium
Read Performance (degraded)	N/A	Medium	High	Low	Low	Low	High	Medium	Medium
Write Performance (degraded)	N/A	High	High	Low	Low	Low	High	Medium	Low
Capacity Utilization	100%	50%	50%	67% - 94%	50% - 88%	50% - 88%	50%	67% - 94%	50% - 88%
Typical Applications	High end workstations, data logging, real-time rendering, very transitory data	Operating system, transaction databases	Operating system, transaction databases	Data warehousing, web serving, archiving	Data warehousing, web serving, archiving	Data archive, backup to disk, high availability solutions, servers with large capacity requirements	Fast databases, application servers	Large databases, file servers, application servers	Data archive, backup to disk, high availability solutions, servers with large capacity requirements

Types of RAID

Types of RAID	Software-Based	Hardware-Based	External Hardware
Description	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. Included in the OS, such as Windows®, Netware®, and Linux. All RAID functions are handled by the host CPU which can severely tax its ability to perform other computations.	Best used for small block applications such as transaction oriented databases and web servers. Processor-intensive RAID operations are off-loaded from the host CPU to enhance performance. Battery-back write back cache can dramatically increase performance without adding risk of data loss.	Connects to the server via a standard controller. RAID functions are performed on a microprocessor located on the external RAID controller independent of the host.
Advantages	Low price Only requires a standard controller	Data protection and performance benefits of RAID More robust fault-tolerant features and increased performance versus software-based RAID	OS independent Build high-capacity storage systems for high-end servers

About Adaptec RAID

Adaptec's RAID technology is a thoroughly tested, proven, and trusted software solution that has been deployed in millions of business-critical installations around the world. As such, our hardened RAID code is the most robust and reliable data protection technology available today. In fact, Adaptec RAID ships inside the most popular name-brand servers — more than 1.5 million servers each year!

adaptec

Adaptec, Inc.
691 South Milpitas Boulevard
Milpitas, California 95035
Tel: (408) 945-8600
Fax: (408) 262-2533

Literature Requests:
US and Canada: 1 (800) 442-7274 or (408) 957-7274
World Wide Web: http://www.adaptec.com
Pre-Sales Support: US and Canada: 1 (800) 442-7274 or (408) 957-7274
Pre-Sales Support: Europe: Tel: (44) 1276-854528 or Fax: (44) 1276-854505

P/N: 666849-011 Printed in USA 10/05 3994_1.11