

SPECIALIST READING

A Find the answers to these questions in the text and table below.

- 1 Give two reasons why server computers often have connected hard drives.
- 2 Why is RAID 0 particularly suited to imaging and scientific work?
- 3 What is the advantage of using drive mirroring?
- 4 To store data, RAID levels higher than 1 require:
 - a At least double the disk space
 - b Up to about a third more disk space
 - c Less than half the disk space
- 5 Where is the backup data stored in a RAID 5 system?
- 6 Which levels of RAID can reconstruct data lost in failed drives from the backup data spread across the remaining drives in the array?
- 7 Which level of RAID is the fastest?

RAIDING HARD DRIVES

Server manufacturers connect hard drives to ensure that data is adequately protected and can be quickly accessed. Computer engineers call such an arrangement a redundant array of inexpensive disks (RAID). By arranging drives in sets, users hope to take advantage of the higher seek times of smaller drives. A special hard disk controller, called a RAID controller, ensures that the RAID array's individual drives are seen by the computer as one large disk drive.

RAID schemes are numbered, with higher numbers indicating more elaborate methods for ensuring data integrity and fault tolerance (or a computer's ability to recover from hardware errors).

B Re-read the text and table to find the answers to these questions.

- 1 Match the terms in Table A with the statements in Table B.

Table A

- | | |
|---|-----------------|
| a | RAID |
| b | RAID controller |
| c | An array |
| d | Striping |
| e | Mirroring |
| f | Parity data |

	Raid 0	Raid 1	Raid 2–4	Raid 5
Fault tolerance?	No	Yes	Yes	Yes
What does it do?	Called disk striping, RAID 0 breaks data into blocks that are spread across all drives rather than filling one before writing to the next.	Called disk mirroring, RAID 1 uses two identical drives: data written to the first is duplicated on the second.	RAID 2–4 are rarely used and simply enhance the striping provided by other RAID levels.	Called striping with parity, the popular RAID 5 writes error-correcting, or parity, data across available drives.
What are the advantages?	Improved disk I/O throughput – the fastest of all RAID configurations as it distributes read/write operations across multiple drives. Good for imaging and scientific work where speed is important.	If either drive fails, the other continues to provide uninterrupted access to data.	2 enhances 0 by using additional drives to store parity data. 3 enhances 2 by requiring only one error-checking drive. 4 builds on 3 by using larger block sizes, boosting performance.	If one drive fails, its contents are recovered by analysing the data on the remaining disks and comparing it with the parity data.
What are the disadvantages?	The failure of any single drive means the entire array is lost.	Inefficient use of disk space.	Uses dedicated disks to store the parity data used to reconstruct drive contents. Up to 30% more hard disk space needed than 1.	Not as fast as RAID 0.

Table B

- i Information which is used to restore data if one of the RAID drives fail
- ii A process of spreading data across a set of disks
- iii Redundant array of inexpensive disks
- iv A set
- v A device for controlling a set of hard disks
- vi The technique of writing the same information to more than one drive

2 Mark the following statements as True or False:

- a Small disks tend to have lower seek times than large disks.
- b RAID controllers make one large hard disk act like a set of small disks.
- c In RAID systems, one disk is filled with data before the next disk is used.
- d A higher numbered RAID array uses a more elaborate system to protect the integrity of data.
- e RAID 0 provides good data recovery.
- f Small file servers do not usually use RAID level 3.