

RAID

REDUNDANT ARRAY OF INDEPENDENT DISKS



WHAT IS RAID

RAID stands for 'redundant array of independent disks' and is a technology that is used to provide **redundancy, performance increases** or both to hard drives which are logically grouped together as a single unit.

RAID LEVELS



There are different RAID levels which provide different goals, these goals include **reliability, availability** and **performance**. Some of the most common RAID levels used are RAID 0, 1, 5 and 6.

RAID 0

RAID 0 uses a technology called **striping** which means data is read and written to each hard drive in the **logical volume**. For example, two hard drives which are two Terabytes in size will logically show as a single 4 Terabyte volume. RAID 0 is excellent for **speed** and **performance** as the read/write speed of each drive is taken advantage of. RAID 0 should not be used where redundancy or availability is required as a single failure in the logical volume means all data will be lost

RAID 1



RAID 1 uses a technology called **mirroring** which means data is written identically to both drives. For example, two hard drives which are two Terabytes in size will logically show as a single 2 Terabyte volume. RAID 1 offers good performance for reading data but slower write speeds than RAID 0 as data must be written to both drives. RAID 1 offers a layer of redundancy and allows for 1 hard drive failure without losing data.

RAID 5



RAID 5 uses **striping and parity** which means data is striped across all drives to increase performance and the system calculates parity information which is also spread across all drives to give a level of resiliency. This means when a hard drive fails, the parity information can be used to rebuild a drive once it has been replaced. RAID 5 requires a minimum of 3 hard drives and is widely used due to its performance and resiliency capabilities.

