

RAID Defined

RAID is an acronym for Redundant Array of Inexpensive Disks, or Redundant Array of Independent Disks.

RAID

RAID was developed to improve hard drive storage reliability and performance. Originally RAID systems consisted of clustering small, "inexpensive" disk drives together into an array so that the array would appear to the system as a single logical drive. During initial testing, it was discovered that an array of drives actually delivers performance exceeding that of single, more expensive hard drives.



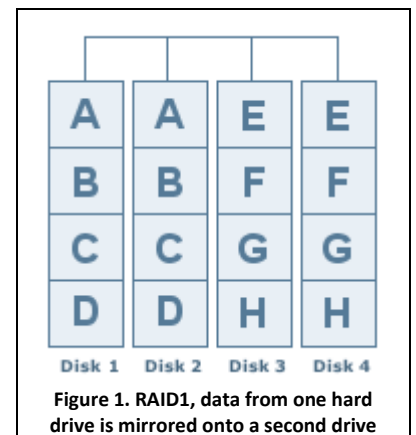
The Mean Time Before Failure (the average time before a failure will occur) in a RAID, was reduced due to the probability of any one drive in the array failing. Consequently, five levels of RAID were originally developed to provide a balance between performance and data protection.

RAID 1

The furnace computer utilizes RAID1 disk mirroring which provides data protection by duplicating all data from a primary drive on a secondary drive. The benefit of this system is it offers the highest data protection.

Since RAID 1 employs the mirroring technique, it offers the advantage of 100% redundancy. If one disk fails, rebuild the lost data from the mirror. RAID 1 requires at least 2 hard drives, and additional hard drives must always be added in pairs. It is ideal for applications where uptime is important and/or data is critical.

In RAID 1 (mirroring without parity or striping), data is written identically to multiple disks (a "mirrored set"). While any number of disks may be used, the furnace system requires only two. The array continues to operate as long as at least one drive is functioning. With appropriate operating system support, there can be increased read performance, and only a minimal write performance reduction. Implementing RAID 1 with a separate controller for each disk in order to perform simultaneous reads (and writes) is sometimes called *multiplexing* (or *duplexing* when there are only 2 disks).



As shown in Figure 1, in a RAID 1 array, data from one hard drive is mirrored onto a second hard drive, so that there are two identical copies of the data. Data from a third drive is mirrored to a fourth drive. RAID 1 drives must always be added in pairs.

Verifying That RAID Is Working on a Dell Optiplex Furnace Computer

Your computer displays information pertaining to your RAID configuration at start-up, before loading the operating system. If RAID is not configured, the message none defined appears under RAID Volumes, followed by a list of the physical drives installed in your computer. If a RAID volume is identified, you can then check the Status field to determine the current state of your RAID configuration. The Status field contains information about the following conditions:

- Normal — Your RAID configuration is functioning properly.
- Degraded — One of your hard drives has failed. The computer is still bootable; however, RAID is not functioning and data is not being copied to the other drive.
- Rebuild — Following a degraded condition, the computer has detected the replacement/connection of a secondary hard drive and will automatically restore the RAID configuration the next time the operating system loads.

RAID Data Recovery

RAID Rebuilding

With their built-in redundancy, RAID systems are able to continue functioning even if a hard drive fails. When this happens however, performance is negatively affected, and the RAID is said to be operating in a degraded, or critical state. This occurs because the lost information must be regenerated "on the fly" from the parity data.

When a failed drive is replaced, the data that was removed from the array with its departure must be regenerated on the new disk by the RAID controller, a process called rebuilding. An array is vulnerable while it is running in a degraded state. Until the failed drive is replaced and its data rebuilt, the array provides no redundancy. Consequently, it's not a good idea to put off rebuilding a degraded array; you're putting your data at risk!

RAID Data Recovery Process

RAID systems are susceptible to the same ailments that plague single hard drives, such as viruses, logical problems, human error, and physical damage. Moreover, due to their complexity, they may suffer from additional points of failure, such as lost server registry configurations, accidental RAID drive reconfigurations, RAID controller failures and multiple drive failures.

If multiple drives fail, or other serious problems occur in a RAID, your data may be compromised.

Under such circumstances, if you fail to make a proper backup, you may have to call a data recovery company. Many have high success rates when it comes to RAID data recovery. Using specialized facilities, equipment, and software, they can even recover data from a moribund hard drive.

Recovering From a Single Hard Drive Failure Using the Intel Matrix Storage Manager on a RAID1 Dell Optiplex Furnace Computer

NOTE: Perform the following steps only after you have replaced the failed hard drive (see the appropriate "Drives" section for your computer). Remember to check the warranty. If the system is still in warranty contact Dell with the Service Tag and/or Express Service Code of the furnace computer. Dell will provide next day on-site installation of the failed drive

1. Turn on or restart your computer.
2. Press <Ctrl><i>< ></i> when you are prompted to enter the Intel RAID Option ROM utility.
3. Under **DEGRADED VOLUME DETECTED**, confirm that the new (non-RAID) drive is listed and then press <Enter>.
4. Under **Disk/Volume Information** confirm that the volume status is *Rebuild*.

NOTE: Volumes with a status of *Rebuild* are rebuilt within the operating system.

5. Use the up- and down-arrow keys to select **Exit**, and then press <Enter>.

Your computer boots to the operating system and begins rebuilding the RAID volume automatically. A dialog box appears and displays the progress of the rebuild.

NOTE: You can use your computer while the computer is rebuilding the RAID level 1 volume.

RAID Levels

Typical features of most common RAID levels are depicted in Table 1.

Table 1. RAID Levels

Level	Description	Min # Disks	Fault Tolerance	Array Failure Rate	Failure Conditions	Benefits/Disadvantages	Image
RAID 0	Disk striping. Block-level striping without parity or mirroring.	2	0 (none)	n	When one drive fails, the entire array is compromised.	Offers Best performance No fault tolerance.	
RAID 1	Disk Mirroring without parity or striping.	2	$n-1$ disks	n	If one drive fails, data is not lost. If both drives fail, the data is lost.	100% redundancy of data/Slower performance and 50% loss of storage space.	
RAID 2	Bit-level striping with dedicated Hamming-code parity.	3	Recover 1 disk failure or repair corrupt data when a corrupted bit's corresponding data & parity are good.	variable	Only one drive may fail and still be recoverable "on the fly".	On the fly data error correction/Extremely high cost.	
RAID 3	Byte-level striping with dedicated parity.	3	1 disk	$n(n-1)r2$	When more than one drive fails, the array is compromised.	High read/write data transfer rates/Complex controller design	
RAID 4	Block-level striping with dedicated parity.	3	1 disk	$n(n-1)r2$	When more than one drive fails, the array is compromised.	High Read/Low Write data transaction rates.	
RAID 5	Block-level striping with distributed parity.	3	1 disk	$n(n-1)r2$	When more than one drive fails, the array is compromised.	High Read data transaction rates/Complex controller design	
RAID 6	Block-level striping with double distributed parity.	4	2 disks	$n(n-1)(n-2)r3$	When more than one drive fails, the array is compromised.		

RAID TERMS

Striping

Striping is a technique which offers the best performance of any RAID configuration. **In a striped array, data is interleaved across all the drives in the array.**

Striping works by splitting up the data and distributing it across multiple drives to increase performance. Performance in a striped array is dependent on the stripe width (the number of drives in the array) and the stripe size (the size of the chunks of data being written across the array).

Striping can occur at two different levels: byte level and block level. Byte level striping involves breaking up the data into bytes and storing them sequentially across the hard drives. Block level striping involves breaking up the data into a given block size. These blocks are then distributed in the same way across the array as in byte level striping.

The stripe size should recommend to bring the most performance from a RAID system depends on the type of application you're using it for. Larger stripes mean fewer accesses to the disk. For this reason, larger stripes are useful for I/O-intensive (Input/Output) applications such as database servers. Smaller stripes on the other hand, mean that data can be accessed more quickly because data chunks are smaller. Consequently, smaller stripes are better suited for throughput-intensive applications such as video production and editing.

Mirroring

Although a striped array may offer the best performance of any RAID configuration, it provides no redundancy. If one drive in the array fails, all of your data will be lost and you may need to consider RAID data recovery options. That's where mirroring comes in. **With mirroring, whatever is written to one drive, gets written simultaneously to another. Thus, there is always have an exact duplicate of your data on the second drive.** This is one of the two data redundancy techniques used in RAID to protect you from data loss. The advantage of this technique is that when one hard drive in the array fails, the system can still continue to operate since there are two copies of the data. Downtime is minimal and rebuilding data from the good copy is relatively easy.

Mirroring also provides a small performance boost over a single non-arrayed drive. Since the mirrored pairs contain the same data, the RAID controller can read data from one drive while simultaneously requesting data from the other. Of course, write speeds will be slower than with other techniques because data must be written twice, once on each drive.

Parity

Parity is an error correction technique commonly used in certain RAID levels. It is used to reconstruct data on a drive that has failed in an array. The RAID controller adds a parity byte to all binary information being written to the array which is an extra byte of data added onto the actual data. **These parity bytes are summed by the controller to equal either an even or an odd number. By analyzing this value, the controller can determine whether the information has been compromised in any way.** To the extent it has, it can replace the data automatically with data from the other drive.

Parity data is created using a logical operation called eXclusive OR (XOR). The controller analyzes the series of 0's and 1's which make up the data, and returns either a TRUE (for even numbers) or FALSE (for odd ones). By using this data, it can "fill in the holes. The XOR logic is used in this way to rebuild corrupted data on the array, thus maintaining integrity.