

# Raid Levels Explained

RAID, or Redundant Array of Independent Disks, integrates multiple disks into a single array to achieve faster performance goals like better hardware failover and improved disk Input/Output performance. We've summarized the primary RAID levels below.

Use our RAID Calculator to see how much raw space will be presented to your operating system (OS) as a storage device based on your desired RAID level. [HostDime servers](#) are fully customizable during the order process to add RAID levels to increase performance.

## RAID Levels Explained

### RAID 0

- Summary: RAID 0 stripes data at the block level across each of the disks in the array. Because the data is striped across multiple disks, you will see improved read and write performance, especially for larger files. This is due to the ability of the array to simultaneously read or write a block to each of the disks in the array. As such, the performance will tend to scale with the number of disks within the array. A RAID 0 can contain two or more disks, with the maximum number being limited based on the number of disks supported by your RAID controller and available drive bays for your server. Because RAID 0 spreads storage blocks across all of the disks in the array, any disk failure will result in the entire array being lost. As such, the risk of data loss due to a drive fault will also scale with the number of disks within the array. We do not recommend using a RAID 0 for production systems.
- Capacity:  $VD = (C * n)$  | VirtualDisk = ( Capacity of one disk \* number of disks )
- I/O Performance: Greatly improved read and write performance
- Fault Tolerance: None, increased risk of failure that scales with number of disks in the array

### RAID 1

- Summary: RAID 1 mirrors data at the block level on two disks. Each time a data block is written to the storage subsystem, your RAID controller will write that same data block on both disks, providing fault tolerance should one of the disks fail. Because of this, the I/O performance characteristics will have slightly lower write performance compared to a single disk, while read performance will be nearly double that of a single disk. This is because when the RAID controller needs to write data it has to do so on both disks, however, when it needs to read data it can simultaneously fetch two different reads from either disk since the data set on both disks is mirrored. Because the data is mirrored on both disks in the array, either one of the disks can fail, as long as one remains functional, without causing the array to fail. NOTE: While most hardware based RAID controllers typically only support two disks in a RAID 1 array, some controllers and software RAID systems may support more than two disks. In these cases, write performance will slightly degrade as the number of disks increase and the read performance will scale with the number of disks.
- Capacity:  $VD = (C * n) / n$  | VirtualDrive = ( Capacity of one disk \* number of disks ) / number of disks
- I/O Performance: Slightly lower write performance with nearly double the read performance
- Fault Tolerance:  $n / n$

### RAID 5

- Summary: RAID 5 stripes data at the block level across each of the disks in the array, while also calculating parity data, which is distributed across the disks in the array to add fault tolerance. With a RAID 5 array, there will be one parity set distributed across the array, which allows for any one disk in the array to fail without causing the array to fail. Because the distributed parity needs to contain parity information for one disk's worth of information, the raw space provided by the array will be the total capacity of all of the disks minus one disk's worth of capacity. This is also why a RAID 5 has a minimum requirement of three disks. The I/O performance for RAID 5 will offer improved read and write performance, however, writes will not scale as linearly as a RAID 0 due to the RAID controller needing to perform the parity calculation and also store the resulting parity data. NOTE: While some software RAID systems support RAID 5, we do not recommend this in a production because all of the parity calculations will need to be done by your server's CPU rather than the dedicated specialized hardware controller that is designed to accelerate these types of calculations. As such, a software RAID 5 will see CPU load scale with disk I/O and can often cause overloads as disk writes increase.
- Capacity:  $VD = (C * n) - C$  | VirtualDrive = ( Capacity of one disk \* number of disks ) - Capacity of one disk
- I/O Performance: Improved read and write performance
- Fault Tolerance:  $n - 1$

### RAID 10

- Summary: RAID 10 is a nested array containing RAID 0 of RAID 1 sets; in other words it is a RAID array of multiple RAID arrays. As such, this RAID level combines some of the performance characteristics of both RAID 0 and RAID 1. The I/O performance will be improved for both read and writes, which will scale differently based on the number of disks in the array. The read performance will scale with the number of disks in the array, while the write performance will scale with the number of RAID 1 sets in the array. For example, a RAID 10 with four disks will have two RAID 1 sets that are then part of a RAID 0, therefore the read performance will be nearly four times a single disk in the array, and write performance will be nearly twice that of a single disk. This RAID level offers a great combination of performance and fault tolerance. Half of the disks in the array can fail without the array failing. This is because one disk in each of the RAID 1 sets can fail without that RAID set failing. Because of these characteristics, you will notice that the minimum number of disks required will be four and the total must be an even number of disks, such that increasing the number of disks in the array is done in pairs of disks. While the number of disks required compared to the raw useable capacity is lower, this RAID level offers the best overall combination of performance and fault tolerance. This is a great option for read heavy databases or mixed work loads.
- Capacity:  $VD = (C * n) / n$  | VirtualDrive = ( Capacity of one disk \* number of disks ) / number of disks

- I/O Performance: Improved write performance with greatly improved read performance
- Fault Tolerance:  $n / 2$

Even with these RAID options, any server is vulnerable to data loss in the case of drive failure. We recommend keeping multiple backup sets including both locally and off-site. Check out our [Disaster Recovery Solutions](#) and our [Remote Backup Service](#) for serious data protection.

If you have any questions, don't hesitate to contact our [support team](#), available around-the-clock.