

RDX® QuikStation® Best Practices with Protected (RAID 5/RAID 6) Volumes

September 2021

Overview

Throughout the storage industry, RAID (Redundant Array of Independent Disks) variants are a standard means of providing more flexibility, redundancy, and security than individual storage devices. RAID 5 and RAID 6, in particular, are relied upon for redundant storage with minimal additional hardware to protect data, and have been for many years.

However, in the intervening decades since RAID 5 technology was developed, the increasing capacity of storage devices and thus a corresponding increase in time required to rebuild a RAID 5 volume after a device failure has made RAID 6 or RAID10 the preeminent choices for redundant storage. RAID 5 has become a poorer option, particularly with large (4 TB or larger) constituent devices.

Description

In brief, the combination of the statistical probability of uncorrectable drive errors and the fact that larger devices require more reading means that RAID 5 in and of itself is not as robust as many people expect. A RAID 6 array the same capacity as a RAID 5 array is better. The primary issue is the total amount of data to be read during reconstruction.

A deep technical analysis of the problem space can be found here:

<https://www.digistor.com.au/the-latest/Whether-RAID-5-is-still-safe-in-2019/>

A (nearly) worst-case scenario calculator can be found at:

<http://magj.github.io/raid-failure/>

We highly recommend you use this calculator to get a rough idea of what works best in your scenario.

NOTE: The Unrecoverable Read Error Rate for most traditional hard drives is 10^{14} , while for most SSDs it is 10^{16} , which is vastly better.

To be realistic, it's still unlikely that an unrecoverable read error would occur in parts of the devices that are in use, and yet more unlikely that such an error would occur in data that is useful/needed (for example, such an error in a JPEG image is likely to be completely unnoticeable, while the same error in a database could be unrecoverable), but it's important for creators and consumers of data to understand the technical risk of RAID 5 devices to be able to make their own risk assessments.

At Overland-Tandberg, we use RAID 5 for a Protected Volume with four RDX devices, and RAID 6 for a Protected Volume with eight RDX devices. We strongly encourage the use of eight-device Protected Volumes when possible (in a QuikStation 8), and advise careful analysis of the risks associated with RAID 5 and large components when selecting Protected Volumes on a QuikStation 4.

When in doubt, smaller cartridges afford better recovery odds than larger ones, and SSD-based cartridges afford significantly better recovery odds than HDD-based cartridges.