LITE-ON Storage is now Solid State Storage Technology Corporation

A Quick Exploration of RAID

Introduction to RAID

RAID, or Redundant Array of Inexpensive Disks, is a way of arranging two or more storage devices in order to improve one or more characteristics in comparison to the devices operating singly. The minimum amount of storage devices, or drives, ranges from two to four, depending on the type of RAID. Different types of RAID provide different benefits with implementation generally seeking to provide redundancy – hence the name – with options to improve reliability. In the balance are other factors, such as capacity and performance. The choice of RAID scheme is largely contingent upon intended workload and overall storage goals.

RAID is a well-understood technology even though improvements have been iterated since its inception. An array of drives function as one, exposed as a single logical device, through the use of striping, mirroring, and/or parity. Utilizing multiple drives in this manner can improve fault tolerance and performance. RAID can still be vulnerable to certain types of failures and it does introduce complexity. There's more than one way to implement and manage RAID. It's possible to utilize both hard drives and solid state drives in a RAID configuration. Typically, though, RAID is chosen to improve storage resiliency at the cost of lower effective capacity.

RAID Configurations

RAID has is origins in JBOD, or Just a Bunch of Disks, which is the combination of multiple drives into a singular array. In that case the disks can be disparate – that is, of various capacities with varying performance – while RAID works best with similar drives. There are multiple kinds of RAID but some schemes are more common and useful than others. For the most part, the typical RAID levels utilized include: RAID-0, RAID-1, RAID-5, RAID-6, and RAID-10.

These all offer some combination of striping, mirroring, and parity, as mentioned above. Mirroring is an obvious implementation – a given pair of drives are effectively clones of each other – while striping involves size-selected blocks on every drive in the RAID. Parity provides fault tolerance through XOR and other types of error correction.

RAID-0 operates with only striping so has no actual redundancy but allows the drives to work together for full capacity and increased performance, especially with queued I/O. Stripe size can vary depending on expected file sizes and device characteristics. RAID-1 is the basic mirror configuration at only one-half capacity but with the ability to recover from a lost drive. RAID-10, or RAID1+0, is a combination of these two, creating a striped set from mirrored drives; this requires at least four drives and offers a balance between redundancy, performance, and capacity.

RAID-5 and RAID-6 utilize distributed parity in single- and double-parity mode so that the failure of any one or two drives, respectively, does not destroy the array. Losing a drive will still negatively impact performance. Basic parity operates with the simple XOR, or exclusive-or, logical operation, which has many other applications in storage. This is sufficient for RAID-5 but RAID-6 requires two distinct parities which alleviates the weaknesses of RAID-5, effectively increasing reliability, particularly with larger arrays.

Types of RAID

A RAID implementation may be software-, hardware-, or firmware-based. Software is the most common and can be done in the operating system, for example through Disk Management or Storage Spaces on Microsoft Windows. Firmware-assisted RAID, or hybrid RAID, is done with the assistance of a motherboard's storage controller. Hardware-based RAID, which is most common with servers and in enterprise, has dedicated hardware for management of the RAID. It's also possible to create informal RAID sets through UEFI.

Software-based RAID has more overhead, i.e. from CPU and memory, is more vulnerable, and can lack portability. However, it is cheaper and relatively easy to implement. Firmware-assisted RAID is similar – portability usually is limited to that specific type of storage controller, for example – but has more flexibility with booting. Hardware-based RAID has a dedicated controller and possibly memory or cache as well, with the option of power loss protection and a write-back mode through battery support. It is more expensive and complex, however, but offers superior configuration through an Option ROM. The abilities to off-load overhead and cache for the RAID have significant benefits.

SSDs and RAID

SSDs share many characteristics with RAID internally. They operate on the principle of parallelization, a type of striping, across multiple channels and flash dies. They utilize XOR logic to scramble data internally, including with encryption. SSDs also rely on a form of RAID – RAIN, or Redundant Array of Independent NAND – to provide parity protection for user data. In this way SSDs are different than HDDs but, further, their performance characteristics lend themselves to a different sort of scalability. Additionally, the interior workings and structure of flash – for example, with typically 16kB physical pages and the need for over-provisioning – mean that RAID may have to be adjusted for SSDs.

SSD array IOPS can actually reach a state where overhead becomes a very real concern. As discussed in many of our other articles, the NVMe™ Specification is designed to address these challenges by improving the management of solid state storage in particular. This includes utilizing the drives themselves to offset overhead in various ways, for example through on-drive compute and PMR. However, in general, NVMe™ allows for superior efficiency and PCIe® connectivity with one primary concept being abstraction – that is, for storage to be represented as logical but divisible assets – which works comfortably with the concept of RAID.

Summary

RAID provides a way to leverage two or more disks in a way that will improve reliability through redundancy. This occurs through the use of mirroring or parity in various configurations, trading capacity for resiliency. RAID can be managed by software, firmware, or hardware, allowing additional flexibility depending on specific storage needs. Larger arrays in particular will benefit from hardware-based RAID-6, for example. Parity implementations rely on XOR or other error correction to ensure array availability – the loss of one or more drives should not be catastrophic. Coupled with expanded connectivity, thanks to PCIe® and the NVMe™ Specification, these resources can be universally and reliably available.

As cloud services continue to expand and become more flexible – the ability to spin up instances and server applications on-the-fly enables unprecedented scalability and efficiency – there is an assorted growth in storage capacity. Performance remains critical to ensure clients are served seamlessly, which requires the low latency of SSDs coupled with 24/7 availability. Backups are part of this scheme but hot-swappable arrays supported through redundancy – RAID – make sure the data world never stops turning.

*All product and company names may be trademarks or registered trademarks of their respective holders.

Our SSD Solutions

PCIe™ - Our ED1 Series is a powerful, high performance SSD made for edge storage applications. It comes in M.2 and U.2 form factors.

SATA - Our ER2 SSD Series delivers affordability and performance with superior random read/write speeds of up to 90,000/45,000 IOPS. It comes in M.2 and 2.5" form factors.

Please contact our Solid State Storage Technology Corp. expert for more information.

*Specifications and features are subject to change without prior notice. Images are samples only, not actual products.

Request Full Specs Sheets

ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

Learn more at www.ssstc.com