Preserving Digital Information in Long-Term Storage

Nick Lee

Jacobs University Bremen, Computer Science
Campus Ring 6, Mailbox 347, Bremen 28759, Germany
ni.lee@jacobs-university.de

With an ever-increasing portion of our lives being encoded and stored as digital information, the need for robust storage systems capable of protecting against physical, software, and random faults is greater than ever. With the widespread availability of high performance hardware, relative to where the consumer industry was five years ago, it is now possible to add more protections and to integrate those mechanisms to an extent not possible before. Through the power of modern technology, it should be possible, given proper physical maintenance, to preserve data with a high degree of confidence for decades at a time.

The most important aspect of data preservation is the protection of the physical media; without a physically sound storage device, data storage, let alone preservation is laughable at best. Today, mass storage typically takes the form of either magnetic hard disks or solid state storage. While both forms of storage are vulnerable to hardware failure, magnetic media is particularly so, and requires either constant backups or hardware redundancy for fault tolerance.

This redundancy often comes in the form of a Redundant Array of Independent Disks (RAID)–unfortunately, while this is a mature technology, it is affected significantly by power disruptions, where the array of pools is rendered inconsistent by what is known as the "Write Hole"[1], and data may be lost–essentially, a recoverable hardware fault has occurred, but the very design of RAID resulted in data loss. Economic considerations must be taken into account as well, as the maximum possible efficiency of the array is $\frac{(n-1)\cdot S}{n\cdot S}\%$ in a RAID5 configuration, which commonly requires four or more drives to operate. RAID 1 also is used as a redundancy scheme, but it is even less efficient at only 50% usable space.

In recent years, software vendors moved to address design issues such as the write hole by introducing new standards such as Oracle's "RAID-Z"[2] and the currently-in-development "btrfs-RAID"[3], which both promise to close the write hole and to provide a better level of fault tolerance. In fact, the the ZFS and btrfs filesystems are particularly well-suited for archival purposes, as they preform cryptographic verification of *all* data read from the drive in an effort to prevent what is known as "bit-rot", the random change of bit states in storage. Through cryptographic verification, effective software, and hardware redundancy, information and backups can be effectively stored without fear of loss or corruption.

In this paper, the evolution of Unix data storage schemes in recent years will be explored, as will some of the technologies utilized.

References

[1] Jeff Bonwick. RAID-Z (Jeff Bonwick's Blog), 2005.

[2] FreeBSD. 20.2. The Z File System (ZFS), 2014.

[3] Jim Salter. Bitrot and atomic COWs: Inside "next-gen" filesystems | Ars Technica, 2014.