

Advantages of OneBlox's Object-based Filesystem



Introduction

Traditional RAID concepts are out of touch with modern storage architectures, placing serious limits on performance, efficiency, and flexibility. Modern storage solutions are moving “post-RAID” with advanced data protection and layout strategies that take advantage of the scale and complexity of today’s datacenter.

Challenges

When RAID was proposed in the 1980’s, it was designed to fit within the limits of available technology. Flash memory had not yet been invented, and applications were seriously limited in terms of storage performance. RAID improved the performance picture and some RAID levels added storage availability as well since the array could continue running after one of the included hard disk drives went off-line. But today’s high-performance storage environment needs more than the moderate levels of performance, availability, and data protection that RAID allows.

RAID was developed as a way to increase performance, since hard disk drives could (and still can) deliver only a hundred or so application I/O operations per second (IOPS). By combining multiple hard disks into a single virtual drive, a RAID set could deliver more IOPS. For example, a striped set (RAID 0) doubles performance, while a parity set (commonly RAID 5) offers somewhat less improvement. Even so, the fastest hard disk RAID configuration cannot reach even ten percent of the performance of an ordinary solid state drive (SSD).

Flexibility was never part of the RAID equation: It permanently “marries” whole identical hard disk drives, with most configurations offering no possibility of expansion as data grows. With RAID, everything is statically configured: Disks are grouped in sets, placed in shelves, behind controllers. Contrast this to the rapidly-shifting modern virtual data center, where virtualized or containerized applications are regularly moved from server to server and even site to site. Modern infrastructure is built of diverse components rather than identical drives.

Another way in which RAID is out of step with modern storage needs is data protection. RAID was designed to survive the failure of a single hard disk drive, and then only in certain configurations. The question of data integrity was left to other levels of the storage stack, since implementing robust error checking and recovery would have required more CPU power than was economically available in the 1980’s. But the rapid expansion of hard disk drive capacity increases the mathematical likelihood of data loss in RAID systems, reaching unacceptable levels with multi-TB hard drives.

Solutions

Modern enterprise and cloud storage solutions either cover up the deficiencies of RAID with layers of software or dispense with it altogether. Two decades ago, the first storage virtualization layers were adopted to overcome the limits of simple RAID sets. These are nearly ubiquitous today, from the array to the hypervisor to the operating system.

But software can only do so much to mask the inefficiency and inflexibility of the underlying RAID architecture, especially as flash memory has been incorporated in storage systems as cache or tier of storage. Advanced storage solutions use data layouts that leverage every component in a system, placing data in chunks across multiple types of storage. In this way, underlying components like hard disks or solid state drives are no longer a limiting factor in system design.

OneBlox Object-based Filesystem

As our engineers were rethinking storage, we at StorageCraft decided that classic RAID was no longer useful. Instead, we developed a unified storage layer spanning every hard disk drive in a ring of OneBlox appliances. This is made possible by our integrated cloud-inspired object store, which treats files as objects in a data store rather than blocks on disk.

When a file is written to the OneBlox ring, it is divided into smaller objects that are dispersed throughout the system according to data protection needs and after inline deduplication and compression. The system is entirely dynamic, and data is re-balanced to match the hardware deployed by our customers. Individual drives and OneBlox appliances can be added and removed as needed, and they don't need to be identical: The system will optimally utilize whatever hardware it finds.

An StorageCraft solution will be more efficient than a typical RAID array as well. Rather than "short stroking" drives to maximize performance or letting some drive sets sit idle due to the limits of RAID allocation, each drive and OneBlox appliance shares an equal part of the storage workload. This high-utilization design increases capacity efficiency and reduces total cost even compared to basic RAID arrays while offering a full set of storage management features. Plus, all data is de-duplicated and compressed before being stored, increasing capacity utilization even further.

OneSystem protects data better than traditional RAID systems as well. Data protection levels always exceed the dual-drive protection of RAID 6 since our widely-dispersed data layout reduces rebuild times. Shorter rebuilds mean less chance for another drive to fail. In addition, all data is protected against corruption with cryptographic checksums, something some RAID systems still do not offer.

Conclusion

StorageCraft OneBlox takes a modern approach to data layout that yields real benefits for our customers. They get better performance, more flexibility, greater capacity utilization, and increased data protection without the high price tag.