RED DATA
An RPI-CS, Inc. division

RPICS

Zoned Storage & blockOS for Software-Defined Storage

Background

During the last several years, Zoned Storage has become more popular due to it usage in hyperscale datacenter environments, enabling high-capacity storage devices while giving the host system more control to optimize for specific storage use cases. Zoned Storage is a standards-based model for managing data on storage devices by dividing their physical address space into distinct zones, with data in each zone written sequentially. The primary benefits are increased capacity (i.e., lower cost per Gigabyte) for high-density recording methods such as Singled Magnetic Recording (SMR) on hard disk drives (HDDs) and Quad-Level Cell (QLC) based NAND Solid State Devices (SSDs).

Use of Zoned Storage - and particularly SMR HDDs[1] – is most impactful in large-scale datacenter deployments where total cost of ownership is key in the decision-making process. In these deployments, traditional RAID for data protection is not common, and data durability is typically handled with replication, erasure coding, or both. These techniques are more SMR-friendly than traditional RAID. Modern software-defined storage typically includes the features needed for enterprise-class uptime and data resiliency, as well as striping data across multiple drives to aggregate throughput and offer better performance than single-drive solutions. In addition, software-defined storage or SDS, offers significant flexibility in the ability to rapidly expand the scale of storage, to mix storage device capacities, and to include multiple storage device types, such as HDDs and SSDs, with different data placement techniques intended to achieve optimal performance and cost goals.

To date, the majority of commercially available SDS products in the market do not natively support Zoned Storage or SMR; some use existing SMR capable filesystem such as btrfs or xfs. However, layering the abstraction of SDS on top of an already-abstracted file system (such as btrfs, xfs) to manage data placement create additional challenges for performance or reliability. Compare that against solutions with native SMR support: they are architected and tested with the specific zoned nature and unique write rules of SMR in mind.

blockOS and Open Archive

RedData's blockOS is a low-overhead object store that manages Zoned Storage devices. It is designed to support host-managed SMR HDDs and Zoned Namespaces SSDs, in addition to managing devices that implement traditional block storage protocols for SATA, SAS, or NVMe. When deployed with RedData's Open Archive, it supports automatic device and volume management, optional inline compression and encryption, erasure coding as well as volume grouping of erasure encoded zones across clusters of storage devices. blockOS can be paired with RedData's software drivers to emulate a tape library, or an API providing get/put object data management services via a web browser or command line interface. When combined with the RedData's ZeroWatt power saving

[1] Shingled Magnetic Recording is a hard disk drive technology that increases data density and storage capacity by allowing data tracks to overlap like shingles on a roof. Unlike conventional Perpendicular Magnetic Recording (PMP) where tracks are written side-by-side, SMR writes a new track that partially overwrites the previous one, enabling more tracks to fit on the platter. This design provides incremental storage capacity per HDD when compared to PMR, thus lowering the cost per gigabyte.

RPI-CS, Inc.
866.938.7775
6747 Katella Avenue, Cypress, CA 90630

technology individual storage devices are powered down as much as possible for optimal power and cost savings with tape-like total-cost of ownership.

RedData's Open Archive is a software-defined storage (SDS) solution that simplifies data lifecycle management. Open Archive extends non-proprietary primary file systems with transparent tiering to automatically move inactive data to lower-cost archive tiers. Additionally, it can stub-out (dehydrate) file system extents leaving only the metadata for very cost-efficient use of tier-0 resources.

Open Archive is scalable by design and supports:

- Tiered storage for data access to distributed systems via a single namespace.
- Uses any available storage media (HDD, SSD, tape, cloud) as archive storage or replication copy.
- Catalog and index services with an extensive and extensible metadata model for searching and locating files from any attached client.
- Data protection for all files with a user-defined policy for file versioning, supporting file or file system recovery to any point-in-time or to another location.
- Policy driven tools for managing data pools, storage media, archive sets, storage capacity, provisioning, and quotas.
- Fine-grained power management at the individual storage device, device groups, or storage enclosure level.

For more information on RedData's Open Archive and blockOS, please contact us: inquiries@rpics.com.